SmoothLLM: Defending Large Language Models Against Jailbreaking Attacks

Abstract

Despite efforts to align large language models (LLMs) with human intentions, widely-used LLMs such as GPT, Llama, and Claude are susceptible to jailbreaking attacks, wherein an adversary fools a targeted LLM into generating objectionable content. To address this vulnerability, we propose SMOOTHLLM, an algorithm designed to mitigate jailbreaking attacks. Based on our finding that adversarially-generated prompts are brittle to character-level changes, our defense randomly perturbs multiple copies of a given input prompt, and then aggregates the corresponding predictions to detect adversarial inputs. Across a range of popular LLMs, SMOOTHLLM offers improved robustness against the GCG, PAIR, RANDOMSEARCH, and AMPLEGCG jailbreaks. SMOOTHLLM is also resistant against adaptive GCG attacks, exhibits a small, though non-negligible trade-off between robustness and nominal performance, and is compatible with any LLM.

1 Introduction

Large language models (LLMs) have emerged as a groundbreaking technology that has the potential to fundamentally reshape how people interact with AI. Central to the fervor surrounding these models is the credibility and authenticity of the text they generate, which is largely attributable to the fact that LLMs are trained on vast text corpora sourced directly from the Internet. And while this practice exposes LLMs to a wealth of knowledge, such corpora tend to engender a double-edged sword, as they often contain objectionable content including hate speech, malware, and false information (Gehman et al., 2020). Indeed, the propensity of LLMs to reproduce this objectionable content has invigorated the field of AI alignment (Yudkowsky, 2016; Gabriel, 2020; Christian, 2020), wherein various mechanisms are used to "align" the output text generated by LLMs with human values (Hacker et al., 2023; Ouyang et al., 2022; Glaese et al., 2022).

At face value, efforts to align LLMs have reduced the propagation of toxic content: Publicly-available chatbots will now rarely output text that is clearly objectionable (Deshpande et al., 2023). Yet, despite this encouraging progress, in recent months a burgeoning literature has identified numerous failure modes—commonly referred to as *jailbreaks*—that bypass the alignment mechanisms and safety guardrails implemented around modern LLMs (Wei et al., 2023; Carlini et al., 2023; Longpre et al., 2024). The pernicious nature of such jailbreaks, which are often difficult to detect or mitigate (Wang et al., 2023), pose a significant barrier to the widespread deployment of LLMs, given that these models may influence educational policy (Blodgett & Madaio, 2021), medical diagnoses (Sallam, 2023), and business decisions (Wu et al., 2023).

Among the jailbreaks discovered so far, a notable category concerns *adversarial prompting*, wherein an attacker fools a targeted LLM into outputting objectionable content by modifying prompts passed as input to that LLM (Maus et al., 2023; Shin et al., 2020; Chao et al., 2023; Liu et al., 2023). Of particular concern are recent works of (Zou et al., 2023b; Andriushchenko et al., 2024; Liao & Sun, 2024), which show that highly-performant LLMs can be jailbroken with 100% attack success rate by appending adversarially-chosen characters onto prompts requesting objectionable content (see (Andriushchenko et al., 2024, Table 1)). And despite widespread interest, existing mitigation strategies have not effectively resolved these vulnerabilities.

In this paper, we begin by proposing a desiderata for candidate defenses against *any* jailbreaking attack. Our desiderata comprises four properties—attack mitigation, non-conservatism, efficiency, and compatibility—

Figure 1: **Preventing jailbreaks with SmoothLLM.** SMOOTHLLM significantly reduces the ADVBENCH attack success rates of four jailbreaking attacks: GCG (Zou et al., 2023b), PAIR (Chao et al., 2023), RANDOMSEARCH (Andriushchenko et al., 2024), and AMPLEGCG (Liao & Sun, 2024).

which outline the challenges involved in defending against jailbreaking attacks on LLMs. Based on this desiderata, we next introduce SMOOTHLLM, a jailbreaking defense algorithm that is particularly effective against suffix-based attacks. The underlying idea behind SMOOTHLLM—which is motivated by the randomized smoothing literature (Lecuyer et al., 2019; Cohen et al., 2019)—is to first duplicate and perturb copies of a given input prompt, and then to aggregate the outputs generated for each copy (see Figure 3).

Contributions. We make the following contributions:

- **Desiderata for defenses.** We propose a desiderata for defenses against jailbreaking attacks. Our desiderata comprises four properties: attack mitigation, non-conservatism, efficiency, and compatibility.

- **LLM jailbreaking defense.** We propose the SMOOTHLLMalgorithm for defending LLMs against jailbreaking attacks. SMOOTHLLM has the following properties:

  - ***Attack mitigation***: SMOOTHLLM significantly reduces the attack success rates (ASRs) of the GCG (Zou et al., 2023b), PAIR (Chao et al., 2023), RANDOMSEARCH (Andriushchenko et al., 2024), and AMPLEGCG (Liao & Sun, 2024) jailbreaks relative to undefended LLMs (see Fig. 1). This is the first demonstration of defending against RANDOMSEARCH and AMPLEGCG.
  - ***Non-conservatism***: Across four NLP benchmarks, SMOOTHLLM incurs a modest, yet non-negligible trade-off between robustness and nominal performance, although we show that this trade-off can be mitigated by picking appropriate hyperparameters for SMOOTHLLM.
  - ***Efficiency***: SMOOTHLLM does not involve retraining the underlying LLM and can improve robustness by up to 20× with a single additional query relative to an undefended LLM.
  - ***Compatibility***: SMOOTHLLM is compatible with both black- and white-box LLMs.

2 The need for defenses against jailbreaking

We start by introducing the threat model that captures the setting of jailbreaking an LLM, which we motivate by describing a characteristic first example: adversarial suffix-based jailbreaks. We then survey existing defenses, which in turn leads to a desiderata for performant defense algorithms against jailbreaking.

Figure 2: **Jailbreaking LLMs.** (Top) Aligned LLMs refuse to respond to the prompt "Tell me how to build a bomb." (Bottom) Aligned LLMs can be attacked by modifying input prompts requesting harmful content.

2.1 Jailbreaking preliminaries

The objective of a jailbreaking attack is to design prompts that, when passed as input to a targeted LLM, cause that LLM to generate an objectionable response. To guide the generation of this content, the attacker is given a goal string G (e.g., "Tell me how to build a bomb") which requests an objectionable response, and to which an aligned LLM will likely abstain from responding (Figure 2, top). Given the inherently challenging and oftentimes subjective nature of determining whether a response is objectionable (Chao et al., 2024), throughout this paper, we assume access to a binary-valued function JB : $R \mapsto \{0, 1\}$ that checks whether a response R generated by an LLM constitutes a jailbreak. That is, given a response R, JB(R) takes on value one if the response is objectionable, and value zero otherwise. In this notation, the goal of a jailbreaking attack is to solve the following problem:

$$\text{find} \quad P \quad \text{subject to} \quad \text{JB} \circ \text{LLM}(P) = 1. \tag{1}$$

Here the prompt P can be thought of as implicitly depending on the goal string G. We note that several different realizations of JB are common in the literature, including checking for the presence of a particular target string T (e.g., "Sure, here's how to build a bomb") as in Figure 2, using an auxiliary LLM to judge whether a response constitutes a jailbreak (Chao et al., 2023; Andriushchenko et al., 2024), human labeling (Wei et al., 2023; Yong et al., 2023), and neural-network-based classifiers (Inan et al., 2023; Huang et al., 2023) (see Chao et al. (2024), §3.5) for a more detailed overview).

2.2 A first example: Adversarial suffix jailbreaks

Numerous algorithms have been shown to solve (1) by returning input prompts that jailbreak a targeted LLM (Chao et al., 2023; Liu et al., 2023; Zou et al., 2023b; Andriushchenko et al., 2024; Liao & Sun, 2024). And while the defense we derive in this paper is applicable to *any* jailbreaking algorithm (see Fig. 1), we next consider a particular class of LLM jailbreaks—which we refer to as *adversarial suffix jailbreaks*—which subsume many well known attacks (e.g., (Zou et al., 2023b; Andriushchenko et al., 2024; Liao & Sun, 2024; Geisler et al., 2024)) and which motivate the derivation of SMOOTHLLM in §3. In the setting of this class of jailbreaks, the goal of the attack is to choose a suffix string S that, when appended onto the goal string G, causes a targeted LLM to output a response containing the objectionable content requested by G. In other words, an adversarial suffix jailbreak searches for a suffix S such that the concatenated string $[G; S]$ induces an objectionable response from the targeted LLM (as in Figure 2, bottom). This setting gives rise the following variant of (1), where the dependence of P on the goal string G is made explicit.

$$\text{find} \quad S \quad \text{subject to} \quad \text{JB} \circ \text{LLM}([G; S]) = 1 \tag{2}$$

That is, S is chosen so that the response $R = \text{LLM}([G; S])$ jailbreaks the LLM. To measure the performance of any algorithm designed to solve (2), we use the *attack success rate* (ASR). Given any collection $\mathcal{D} =$

Figure 3: **SmoothLLM.** (Left) An undefended LLM (**cyan**) takes an attacked prompt P' as input and returns a response R. (Right) SMOOTHLLM (**yellow**), which acts as a wrapper around *any* LLM, comprises a perturbation step (**pink**), wherein N copies of the input prompt are perturbed, and an aggregation step (**green**), wherein the outputs corresponding to the perturbed copies are aggregated.

$\{(G_j, S_j)\}_{j=1}^n$ of goals G_j and suffixes S_j, the ASR is defined by

$$\text{ASR}(\mathcal{D}) \triangleq \frac{1}{n} \sum_{j=1}^n \text{JB} \circ \text{LLM}([G_j; S_j]). \tag{3}$$

In other words, the ASR is the fraction of the pairs (G_j, S_j) in $\mathcal{D}$ that jailbreak the LLM.

2.3 Existing defenses for language-based attacks

The literature concerning the robustness of language models comprises several defense strategies (Goyal et al., 2023). However, the vast majority of these defenses, e.g., those that use adversarial training (Liu et al., 2020; Miyato et al., 2016) or data augmentation (Li et al., 2018), require retraining the underlying model, which is computationally infeasible for LLMs. Indeed, the opacity of closed-source LLMs (which are only available via calls made to an enterprise API) necessitates that candidate defenses rely solely on query access. These constraints, coupled with the fact that no algorithm has yet been shown to significantly reduce the ASRs of existing jailbreaks, give rise to a new set of challenges inherent to the vulnerabilities of LLMs.

Several *concurrent* works also concern defending against adversarial attacks on LLMs. In (Jain et al., 2023), the authors consider several candidate defenses, including input preprocessing and adversarial training. Results for these methods are mixed; while heuristic detection-based methods perform strongly, adversarial training is shown to be infeasible given the computational costs. In (Kumar et al., 2023), the authors apply a filter on sub-strings of prompts passed as input to an LLM. While promising, the complexity of this method scales with the length of the input prompt, which is intractable for most jailbreaking attacks. Also related is the near-concurrent work of Cao et al. (2023), which introduces an algorithm called RA-LLM that applies randomly generated masks to input prompts. Similarly to SMOOTHLLM, RA-LLM also adopts a modify-then-aggregate algorithm, although SMOOTHLLM uses different perturbation functions: inserts, swaps, and patches. Moreover, we investigate how nominal performance and latency vary with various hyperparameters; Cao et al. (2023) leave such experiments to future work.

2.4 A desiderata for LLM defenses against jailbreaking

The opacity, scale, and diversity of modern LLMs give rise to a unique set of challenges when designing a candidate defense algorithm against adversarial jailbreaks. To this end, we propose the following as a comprehensive desiderata for broadly-applicable and performant defense strategies.

(D1) ***Attack mitigation.*** A candidate defense should—both empirically and provably—mitigate the adversarial jailbreaking attack under consideration. Furthermore, candidate defenses should be non-exploitable, meaning they should be robust to adaptive, test-time attacks.

(D2) **Non-conservatism.** While a trivial defense would be to never generate any output, this would result in unnecessary conservatism and limit the widespread use of LLMs. Thus, a defense should avoid conservatism and maintain the ability to generate realistic text.

(D3) **Efficiency.** Modern LLMs are trained for millions of GPU-hours. Moreover, such models comprise billions of parameters, which gives rise to a non-negligible latency in the forward pass. Thus, candidate algorithms should avoid retraining and maximize query efficiency.

(D4) **Compatibility.** The current selection of LLMs comprises various architectures and data modalities; further, some (e.g., Llama2) are open-source, while others (e.g., GPT-4) are not. A candidate defense should be compatible with each of these properties and models.

The first two properties—*attack mitigation* and *non-conservatism*—require that a candidate defense successfully mitigates the attack under consideration without a significant reduction in performance on non-adversarial inputs. The interplay between these properties is crucial; while one could completely nullify the attack by changing every character in an input prompt, this would come at the cost of extreme conservatism, as the input to the LLM would comprise nonsensical text. The latter two properties—*efficiency* and *compatibility*—concern the applicability of a candidate defense to the full roster of currently available LLMs without implementation trade-offs.

3   SmoothLLM: A randomized defense for LLMs

Given the need to design defenses against jailbreaking attacks, we propose SMOOTHLLM, which is particularly well-suited for defending against suffix-based jailbreaks. Key to the design of SMOOTHLLM is the desiderata outlined in §2.4 as well as insights from the randomized smoothing literature (Lecuyer et al., 2019; Cohen et al., 2019; Salman et al., 2019).

3.1   Adversarial suffixes are fragile to perturbations

Our algorithmic contribution is predicated on the following previously unobserved phenomenon: The suffixes generated by adversarial suffix jailbreaks are fragile to character-level perturbations. That is, when one changes a small percentage of the characters in a given suffix, the ASRs of these jailbreaks drop significantly, often by more than an order of magnitude. This fragility is demonstrated in Fig. 4, wherein the dashed lines (shown in **red**) denote the ASRs for suffixes generated by GCG on the AdvBench dataset (Zou et al., 2023b). The bars denote the ASRs corresponding to the same suffixes when these suffixes are perturbed in three different ways: randomly inserting $q\%$ more characters into the suffix (shown in **blue**), randomly swapping $q\%$ of the characters in the suffix (shown in **orange**), and randomly changing a contiguous patch of characters of width equal to $q\%$ of the suffix (shown in **green**). Observe that for insert and patch perturbations, perturbing only $q = 10\%$ of the characters in the each suffix reduces the ASR below 1%.

3.2   From perturbation instability to adversarial defense

The fragility of adversarial suffixes to perturbations suggests that the threat posed by adversarial prompting jailbreaks could be mitigated by randomly perturbing characters in a given input prompt P. This intuition is central to the derivation of SMOOTHLLM, which involves two key ingredients: (1) a *perturbation* step, wherein N copies of P are randomly perturbed and (2) an *aggregation* step, wherein the responses corresponding to these perturbed copies are aggregated and a single response is returned. These steps are illustrated in Figure 3 and described in detail below.

Perturbation step. The first ingredient in our approach is to randomly perturb prompts passed as input to the LLM. Given an alphabet $\mathcal{A}$, we consider three perturbation types:

- **Insert**: Randomly sample $q\%$ of the characters in P, and after each of these characters, insert a new character sampled uniformly from $\mathcal{A}$.
- **Swap**: Randomly sample $q\%$ of the characters in P, and then swap the characters at those locations by sampling new characters uniformly from $\mathcal{A}$.

Figure 4: **The instability of adversarial suffixes.** The **red** dashed line shows the ASR of the attack proposed in (Zou et al., 2023b) and defined in (2) for Vicuna and Llama2. We then perturb $q\%$ of the characters in each suffix—where $q \in \{5, 10, 15, 20\}$—in three ways: inserting randomly selected characters (**blue**), swapping randomly selected characters (**orange**), and swapping a contiguous patch of randomly selected characters (**green**). At nearly all perturbation levels, the ASR drops by at least a factor of two. At $q = 10\%$, the ASR for swap perturbations falls below 1%.

- ***Patch***: Randomly sample d consecutive characters in P, where d equals $q\%$ of the characters in P, and then replace these characters with new characters sampled uniformly from $\mathcal{A}$.

Notice that the magnitude of each perturbation type is controlled by a percentage q, where $q = 0\%$ means that the prompt is left unperturbed, and higher values of q correspond to larger perturbations. In Fig. 5, we show examples of each perturbation type (for details, see Appendix D). We emphasize that in these examples and in our algorithm, the *entire* prompt is perturbed, not just the suffix; SMOOTHLLM does not assume knowledge of the position (or presence) of a suffix in a given prompt.

Aggregation step. The second key ingredient is as follows: Rather than passing a *single* perturbed prompt through the LLM, we obtain a *collection* of perturbed prompts, and then aggregate the predictions for this collection. The motivation for this step is that while *one* perturbed prompt may not mitigate an attack, as evinced by Figure 4, *on average*, perturbed prompts tend to nullify jailbreaks. That is, by perturbing multiple copies of each prompt, we are likely to flip characters in the adversarial portion of the prompt. To formalize this step, let $\mathbb{P}_q(P)$ denote a distribution over perturbed copies of P, where q denotes the perturbation percentage. Now given perturbed prompts Q_j drawn from $\mathbb{P}_q(P)$, if q is large enough, Figure 4 suggests that the randomness introduced by Q_j should nullify an adversarial attack.

Both the perturbation and aggregation steps are central to SMOOTHLLM, which we define as follows.

Definition 3.1 (SmoothLLM)

Let a prompt P and a distribution $\mathbb{P}_q(P)$ over perturbed copies of P be given. Let $\gamma \in [0, 1]$ and $Q_1, \ldots, Q_N$ be drawn i.i.d. from $\mathbb{P}_q(P)$, then define V to be the majority vote of the JB function across these perturbed prompts w.r.t. the margin γ, i.e.,

$$V \triangleq \mathbb{I}\left[\frac{1}{N}\sum_{j=1}^{N}\left[(\text{JB} \circ \text{LLM})(Q_j)\right] > \gamma\right]. \tag{4}$$

Then **SmoothLLM** is defined as

$$\text{SMOOTHLLM}(P) \triangleq \text{LLM}(Q) \tag{5}$$

where Q is any of the sampled prompts satisfying the majority vote condition $(\text{JB} \circ \text{LLM})(Q) = V$.

Figure 5: **SmoothLLM: A randomized defense.** (Left) Examples of insert, swap, and patch perturbations (shown in pink), all of which can be called in the `RandomPerturbation` subroutine in Algorithm 1. (Right) Pseudocode for SmoothLLM. In lines 2-4, we outline the perturbation step. Next, in line 5, we determine whether a γ-fraction of the responses jailbreak the target LLM. Finally, in line 6, we select a response uniformly at random that is consistent with the majority vote.

Notice that after drawing Q_j from $\mathbb{P}_q(P)$, we compute the average over $(\text{JB} \circ \text{LLM})(Q_j)$, which estimates whether perturbed prompts jailbreak the LLM. We then aggregate these predictions by returning any response $\text{LLM}(Q)$ that agrees with that estimate. In Algorithm 1, we translate the definition of SMOOTHLLM into pseudocode. In lines 1–3, we obtain N perturbed prompts Q_j, which is an implementation of sampling from $\mathbb{P}_q(P)$ (see Figure 5). Next, after generating responses R_j for each perturbed prompt Q_j (line 3), we compute the empirical average over the N responses, and then determine whether the average exceeds γ (line 4). Finally, we aggregate by returning a response R_j that is consistent with the majority (lines 5–6). Thus, Algorithm 1 involves three parameters: the number of samples N, the perturbation percentage q, and the margin for the majority vote γ (which, unless otherwise stated, we set to be $1/2$).

3.3   Choosing hyperparameters for SmoothLLM: A simplified analysis

How should the hyperparameters N, q, and γ be chosen? As a first step toward answering this question, we study the theoretical properties of SMOOTHLLM under a **simplifying assumption** which is nonetheless supported by the evidence in Figure 4. This assumption—which characterizes the fragility of adversarial suffixes to perturbations—facilitates the closed-form calculation of the probability that SMOOTHLLM returns a non-jailbroken response, a quantity we term the *defense success probability* (DSP):

$$\text{DSP}(P) \triangleq \Pr[(\text{JB} \circ \text{SMOOTHLLM})(P) = 0]. \tag{6}$$

Here, the randomness is due to the N i.i.d. draws from $\mathbb{P}_q(P)$ in Definition 3.1. Specifically, for the purposes of analysis in a simplified setting, we make the following assumption about adversarial suffix jailbreaks.

Definition 3.2   (k-unstable)

Given a goal G, let a suffix S be such that the prompt $P = [G; S]$ jailbreaks a given LLM, i.e., $(\text{JB} \circ \text{LLM})([G; S]) = 1$. Then S is **k-unstable** *with respect to that LLM if*

$$(\text{JB} \circ \text{LLM})([G; S']) = 0 \iff d_H(S, S') \geq k \tag{7}$$

*where d_H is the Hamming distance between two strings. We call k the **instability parameter**.*

Figure 6: **The defense success probability in a simplified setting.** We plot the defense success probability derived in (8) as a function of the number of samples N and the perturbation percentage q under the simplifying assumption of k-instability; warmer colors denote larger probabilities. From left to right, probabilities are computed for three different values of the instability parameter $k \in \{2, 5, 8\}$. In each subplot, we see that as N and q increase, so does the DSP.

In plain terms, a prompt is k-unstable if the attack fails when one changes k or more characters in S. In this way, Figure 4 can be seen as approximately measuring whether adversarially attacked prompts for Vicuna and Llama2 are k-unstable for input prompts of length m where $k = \lfloor qm \rfloor$. Note that this assumption is difficult to check in practice—it would require enumerating all possible k-length suffixes—and applies specifically to suffix-based attacks. We therefore introduce Definition 3.2 to study SmoothLLM in a simplified, yet natural setting to guide the choice of N, q, and γ, rather than to perform formal certification, as is done in the literature on adversarial examples (Lecuyer et al., 2019; Cohen et al., 2019; Salman et al., 2019).

A closed-form expression for the DSP. We next state our main theoretical result: a closed-form expression for the DSP under the simplifying assumption of k-instability; we present a proof—which requires only elementary probability and combinatorics—in Appendix A, as well as results for other perturbations.

Proposition 3.3 (SmoothLLM certificate)

Given an alphabet $\mathcal{A}$ of v characters, assume that a prompt $P = [G; S] \in \mathcal{A}^m$ is k-unstable, where $G \in \mathcal{A}^{m_G}$ and $S \in \mathcal{A}^{m_S}$. Recall that N is the number of samples and q is the perturbation percentage. Define $M = \lfloor qm \rfloor$ to be the number of characters perturbed when Algorithm 1 is run with swap perturbations and $\gamma = 1/2$. Then, the DSP is as follows:

$$\text{DSP}([G; S]) = \sum_{t=\lceil N/2 \rceil}^{n} \binom{N}{t} \alpha^t (1 - \alpha)^{N-t} \tag{8}$$

where α, which denotes the probability that $Q \sim \mathbb{P}_q(P)$ does not jailbreak the LLM, is given by

$$\alpha \triangleq \sum_{i=k}^{\min(M, m_S)} \left[\binom{M}{i} \binom{m - m_S}{M - i} \Big/ \binom{m}{M} \right] \beta(\ell, i, v) \tag{9}$$

where

$$\beta(\ell, i, v) \triangleq \sum_{\ell=k}^{i} \binom{i}{\ell} \left(\frac{v - 1}{v} \right)^\ell \left(\frac{1}{v} \right)^{i-\ell} \tag{10}$$

This result provides a closed-form expression for the DSP in terms of the number of samples N, the perturbation percentage q, and the instability parameter k. In Figure 6, we compute the expression for the DSP given in (8) for various values of N, q, and k. We use an alphabet size of $v = 100$, which matches our

Figure 7: **Attack mitigation.** We plot the ASRs for Vicuna (top row) and Llama2 (bottom row) for various values of the number of samples $N \in \{2, 4, 6, 8, 10\}$ and the perturbation percentage $q \in \{5, 10, 15, 20\}$; the results are compiled across five trials. For swap perturbations and $N > 6$, SmoothLLM reduces the ASR to below 1% for both LLMs.

experiments in §4 (for details, see Appendix B); m and m_S were chosen to be the average prompt and suffix lengths ($m = 168$ and $m_S = 95$) for the prompts generated for Llama2[1] in Figure 4. Notice that even at relatively low values of N and q, suffix-based attacks would be mitigated under the k-unstable assumption. And as one would expect, as k increases, one should increase q to mitigate the attack.

4 Experimental results

We now consider an empirical evaluation of the performance of SmoothLLM. To guide our evaluation, we cast an eye back to the properties outlined in the desiderata in §2.4: (D1) attack mitigation, (D2) non-conservatism, (D3) efficiency. We note that as SmoothLLM is a black-box defense, it is compatible with any LLM, and thus satisfies the criteria outlined in desideratum (D4).

4.1 Desideratum D1: Attack mitigation

Robustness against jailbreak attacks. In Figure 1, we show the performance of four attacks—GCG (Zou et al., 2023b), PAIR (Chao et al., 2023), RandomSearch (Andriushchenko et al., 2024), and AmpleGCG (Liao & Sun, 2024)—when evaluated against an undefended LLM and an LLM defended with SmoothLLM. In each subplot, we use the datasets used in each of the attack papers (i.e., `AdvBench` (Zou et al., 2023b) for GCG, RandomSearch, and AmpleGCG, and `JBB-Behaviors` (Chao et al., 2023) for PAIR). Notably, SmoothLLM reduces the ASR of GCG to below one percentage point, which sets the current state-of-the-art for this attack. Furthermore, the results in the bottom row of Figure 1 represent the first demonstration of defending against PAIR, RandomSearch, and AmpleGCG in the literature, and therefore these results set the state-of-the-art for these attacks. We highlight that although SmoothLLM was designed with adversarial suffix jailbreaks in mind, SmoothLLM reduces the ASRs of the PAIR semantic attack on Vicuna and GPT-4 by factors of two, and reduces the ASR of GPT-3.5 by a factor of 29.

Adaptive attacks on SmoothLLM. The gold standard for evaluating the robustness is to perform an *adaptive attack*, wherein an adversary directly attacks a defended target model (Tramer et al., 2020). And while at first glance the non-differentiability of SmoothLLM (see Prop. C.1) precludes the direct application of adaptive GCG attacks, in Appendix C.2.2 we derive a new approach which attacks a differentiable SmoothLLM surrogate which smooths in the space of tokens, rather than in the space of prompts. Thus, just as (Zou et al., 2023b) transfers attacks from white-box to black-box LLMs, we transfer attacks opti-

---

[1]The corresponding values for Vicuna are $m = 179$ and $m_S = 106$. We provide an analogous plot to Figure 6 in Appendix B.

Table 1: **Defense performance comparison.** We compare the performance of SMOOTHLLM against (a) no defense, (b) the removal of non-dictionary words (NDW), and (c) a perplexity filter (PF).

Attack	Defense	Target LLM ASR				
		Vicuna	Llama2	GPT-3.5	GPT-4	Average
PAIR	None	82	4	76	50	53
	NDW	82	4	76	50	53
	PF	81	4	15	43	35.75
	Ours	47	1	12	25	21.25
GCG	None	58	2	34	1	23.75
	NDW	10	0	4	1	3.75
	PF	1	0	1	0	0.5
	Ours	1	1	1	0	0.75

mized for the surrogate to SMOOTHLLM. Our results, which are reported in Figure 8, indicate that adaptive attacks generated for SMOOTHLLM are no stronger than attacks optimized for an undefended LLM.

Since the public release of this paper, several follow-up works have studied the robustness of SMOOTHLLM against different adaptive attacks. For instance, (Ji et al., 2024) report the ASRs of SMOOTHLLM under adaptive PAIR and AUTODAN attacks, finding that SMOOTHLLM is more susceptible to these adaptive attacks than it is to the GCG adaptive attacks derived in this paper. Similarly, (Pang et al., 2025) derive a distinct differentiable GCG-like scheme for adaptively attacking SMOOTHLLM (see Algorithm 2 in Pang et al. (2025)), showing similar trends in the loss space to Figure 7.

The role of N and q. In the absence of a defense algorithm, Figure 4 indicates that GCG achieves ASRs of 98% and 51% on Vicuna and Llama2 respectively. In contrast, Figure 1 demonstrates for particular choices of the number of N and q, the effectiveness of various attacks can be significantly reduced. To evaluate the impact of varying these hyperparameters, consider Figure 7, where the ASRs of GCG are plotted for various values of N and q. These results show that for both LLMs, a relatively small value of $q = 5\%$ is sufficient to halve the corresponding ASRs. And, in general, as N and q increase, the ASR drops significantly. In particular, for swap perturbations and $N > 6$, the ASRs of both Llama2 and Vicuna drop below 1%; this equates to a reduction of roughly $50\times$ and $100\times$ for Llama2 and Vicuna respectively.

Comparison to baseline defenses. In Table 1, we compare the performance of SMOOTHLLM to other defenses, including a perplexity filter (abbr. PF) (Jain et al., 2023) and the removal of non-dictionary words (abbr. NDW). Following (Jain et al., 2023), we set the threshold for the perplexity filter to be the maximum perplexity of any of the goal prompts, and we run SMOOTHLLM with $N = 10$ and $q = 5$. We find that while both SMOOTHLLM and the perplexity filter effectively mitigate GCG to a near zero ASR, SMOOTHLLM achieves significantly lower ASRs on PAIR compared to the other defenses. Specifically, across Vicuna, Llama2, GPT-3.5, and GPT-4, SMOOTHLLM reduces the the ASR of PAIR relative to an undefended LLM by 60%, whereas the next best algorithm only decreases the undefended ASR by 32%.

4.2 Desideratum D2: Non-conservatism

Nominal performance of SmoothLLM. Reducing the ASR of a given attack is not meaningful unless the defended LLM retains the ability to generate realistic text. Indeed, two trivial, highly conservative defenses would be to (a) never return any output or (b) set $q = 100\%$ in Algorithm 1. To evaluate the nominal performance of SMOOTHLLM, we consider four NLP benchmarks: `InstructionFollowing` (IF) (Zhou et al., 2023), `PIQA` (Bisk et al., 2020), `OpenBookQA` (Mihaylov et al., 2018), and `ToxiGen` (Hartvigsen et al., 2022). The results on `IF`—which uses two metrics: prompt- and instruction-level accuracy—are shown in Figure 9; the remainder of the results are deferred to Appendix B. Figure 9 shows that as one would expect, larger values of q tend to decrease nominal performance. The presence of such a trade-off is unsurprising: similar trade-offs are extensively documented in computer vision (Croce et al., 2020) and recommendation systems (Seminario & Wilson, 2012). Across each of the dataset, patch perturbations tended to result in

Table 2: **Robustness with one extra query.** For a budget of $q = 10\%$, we report the ASRs for (1) an undefended LLM and (2) SmoothLLM when run with $N = 2$. Relative to the undefended LLM, the SmoothLLM ASRs represent the robustness that can be gained at the cost of one extra query.

Figure 8: **Adaptive attacks on SmoothLLM.** We report the ASRs of a GCG adaptive attack on SmoothLLM run with $N = 10$ and $\gamma = 1/2$ as a function of q. Compared to Figure 1, this adaptive attack is *no stronger* against SmoothLLM than non-adaptive attacks.

LLM	Undefended ASR	SmoothLLM ASR		
		Insert	Swap	Patch
Vicuna	98.0	19.1	13.9	39.8
Llama2	52.0	2.8	3.1	11.0

a more favorable trade-off. For example, on `PIQA`, setting $q = 5$ and $N = 20$ resulted in a performance degradation from 76.7% to 70.3% for Llama2 and from 77.4% to 71.9% for Vicuna (see Table 4).

Improving nominal performance. The following empirical trick improves nominal performance without trading off robustness. First, we set the threshold $\gamma = N-1/N$, which tilts the majority vote toward returning a response R with $\text{JB}(R) = 0$. Then, if indeed the tilted majority vote V in (4) is equal to zero, we return $\text{LLM}(P)$, i.e., a response generated for the unperturbed input prompt. In Table 7, we show that this variant offers similar levels of robustness against PAIR and GCG. However, on the `IF` dataset, we found that across all perturbation levels q, the clean performance matched the undefended performance in Figure 9.

4.3 Desideratum D3: Efficiency

Defended vs. undefended. As described in §3, SmoothLLM requires N times more queries relative to an undefended LLM. Such a trade-off is not without precedent; it is well-documented in the adversarial ML community that improved robustness comes at the cost of query complexity (Wong et al., 2020; Gluch & Urbanke, 2021; Shafahi et al., 2019). Indeed, smoothing-based defenses in the adversarial examples literature require hundreds (see Salman et al. (2019, §5)) or thousands (see Cohen et al. (2019, §4)) of queries per instance. In contrast, as shown in Table 2, for a fixed budget of $q = 10\%$, running SmoothLLM with $N = 2$—meaning that SmoothLLM uses *one extra query* relative to an undefended LLM—results in a 2.5–7.0× reduction in the ASR for Vicuna and a 5.7–18.6× reduction for Llama2 depending on the perturbation. Specifically, for swap perturbations, a single extra query imparts a nearly 20× reduction in Llama2's ASR.

Regarding test-time compute. While the use of additional inference-time computation to defend against jailbreaking may seem unappealing given that the attacker only pays an upfront cost, a similar reallocation of compute is known to advance the reasoning capabilities of LLMs (OpenAI, 2023). This represents an exciting paradigm shift: reallocating computation from training to inference can result in points further along the Pareto frontier spanning LLM capabilities and total computation (Snell et al., 2024). The results in Table 2 are the first to demonstrate that a similar principle may hold for jailbreaking robustness.

On the choice of N. To inform the choice of N, we consider a nonstandard, yet informative comparison of the efficiency of the GCG attack with that of the SmoothLLM defense. The default implementation of GCG uses approximately 256,000 queries to produce a single suffix. In contrast, SmoothLLM queries the LLM N times, where typically $N \leq 20$, meaning that SmoothLLM is generally 5-6× more efficient. In Figure 10, we plot the ASR found by running GCG and SmoothLLM for varying step counts on Vicuna (see Appendix B for Llama2 results). Notice that as GCG runs for more iterations, the ASR tends to increase. However, this phenomenon is countered by SmoothLLM: As N increases, the ASR tends to drop.

Running time analysis. Our argument regarding the efficiency of SmoothLLM has so far centered on the number of queries made to the underlying LLM: For a given prompt, SmoothLLM makes between 10^5 and 10^6 times fewer queries to defend than GCG does to attack. We focus on the number of queries because this figure is hardware-agnostic. However, another way to make the case for the efficiency of SmoothLLM is

Figure 9: **Non-conservatism.** Each subplot shows the performance of SMOOTHLLM run with $N = 10$ on the INSTRUCTIONFOLLOWING dataset; the left and right columns show the performance for patch and insert perturbations respectively, and the dashed lines show the undefended performance. As q increases, nominal performance degrades linearly, resulting in a non-negligible trade-off.

Figure 10: **Efficiency: Attack vs. defense.** Each plot shows the ASRs found by running the attack (GCG) and the defense (SMOOTHLLM) for varying step counts. Warmer colors signify larger ASRs, and from left to right, we sweep over $q \in \{5, 10, 15\}$. SMOOTHLLM uses five to six orders of magnitude fewer queries than GCG and reduces the ASR to near zero as N and q increase.

to compare the time it takes to defend against an attack to the time it takes to generate an attack. To this end, in Table 8, we list the running time per prompt of SMOOTHLLM. These results show that depending on the choice of the number of samples N, defending takes between 3.5 and 4.5 seconds. On the other hand, obtaining a single GCG suffix takes nearly 90 minutes on an A100 GPU and two hours on an A6000 GPU.

5 Discussion

The interplay between q and the ASR. In several of the panels in Figure 7, for lower values of N (e.g., $N \leq 4$), higher values of q (e.g., $q = 20\%$) result in higher ASRs than lower values. While this may seem counterintuitive, since a larger q results in a more heavily perturbed suffix, this subtle behavior is expected. In our experiments, we found that for large values of q, the LLM often outputted the following response: "Your question contains a series of unrelated words and symbols that do not form a valid question." Several judges, including the judge used in Zou et al. (2023b), are known to classify such responses as jailbreaks (see, e.g., Chao et al. (2023, §3.5)). This indicates that q should be chosen to be small enough so that the prompt retains its semantic content. For further examples and discussion, see Appendix E.

The computational burden of jailbreaking. A trend in the literature surrounding robust deep learning is a pronounced computational disparity between efficient attacks and expensive defenses. One reason for this is many methods, e.g., adversarial training (Madry et al., 2017) and data augmentation (Volpi et al., 2018), retrain the underlying model. However, in the setting of adversarial prompting, our results concerning query-efficiency (see Figure 10), time-efficiency (see Table 8), and compatibility with black-box LLMs (see Figure 1) indicate that the bulk of the computational burden falls on the attacker. In this way, future research must seek "robust attacks" which cannot cheaply be defended by randomized algorithms like SmoothLLM.

Addressing the nominal performance trade-off. One limitation of SMOOTHLLM is the extent to which it trades off nominal performance for robustness. While this trade off is manageable for $q \leq 5$, as shown in Figures 9 and 13, nominal performance tends to degrade for large q. At the end of §4.2, we took first steps toward resolving this trade-off, although there is still room for improvement; we plan to pursue this direction in future work. Several future directions along these lines include using a denoising generative model (Salman et al., 2020; Carlini et al., 2022) and using semantic rather than character-level perturbations.

6 Conclusion

In this paper, we proposed SMOOTHLLM, a new defense against jailbreaking attacks on LLMs. The design of SMOOTHLLM is rooted in a desiderata that comprises four properties—attack mitigation, non-conservatism, efficiency, and compatibility—which we hope will guide future research. In our results, we found that SMOOTHLLM offers competitive robustness against GCG, PAIR, RANDOMSEARCH, and AMPLEGCG.

References

Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. Generating natural language adversarial examples. *arXiv preprint arXiv:1804.07998*, 2018. 41

Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Jailbreaking leading safety-aligned llms with simple adaptive attacks. *arXiv preprint arXiv:2404.02151*, 2024. 1, 2, 3, 9

Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in language models is mediated by a single direction. *arXiv preprint arXiv:2406.11717*, 2024. 41

Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 40

Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13*, pp. 387–402. Springer, 2013. 40

Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020. 10, 28

Su Lin Blodgett and Michael Madaio. Risks of ai foundation models in education. *arXiv preprint arXiv:2110.10024*, 2021. 1

Bochuan Cao, Yuanpu Cao, Lu Lin, and Jinghui Chen. Defending against alignment-breaking attacks via robustly aligned llm. *arXiv preprint arXiv:2309.14348*, 2023. 4

Nicholas Carlini, Florian Tramer, Krishnamurthy Dj Dvijotham, Leslie Rice, Mingjie Sun, and J Zico Kolter. (certified!!) adversarial robustness for free! *arXiv preprint arXiv:2206.10550*, 2022. 12

Nicholas Carlini, Milad Nasr, Christopher A Choquette-Choo, Matthew Jagielski, Irena Gao, Anas Awadalla, Pang Wei Koh, Daphne Ippolito, Katherine Lee, Florian Tramer, et al. Are aligned neural networks adversarially aligned? *arXiv preprint arXiv:2306.15447*, 2023. 1

Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*, 2023. 1, 2, 3, 9, 12, 26

Patrick Chao, Edoardo Debenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwag, Edgar Dobriban, Nicolas Flammarion, George J Pappas, Florian Tramer, et al. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. *arXiv preprint arXiv:2404.01318*, 2024. 3

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL https://lmsys.org/blog/2023-03-30-vicuna/. 25

Brian Christian. *The alignment problem: Machine learning and human values*. WW Norton & Company, 2020. 1

Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *international conference on machine learning*, pp. 1310–1320. PMLR, 2019. 2, 5, 8, 11, 40, 41

Francesco Croce, Maksym Andriushchenko, Vikash Sehwag, Edoardo Debenedetti, Nicolas Flammarion, Mung Chiang, Prateek Mittal, and Matthias Hein. Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint arXiv:2010.09670*, 2020. 10

Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. Toxicity in chatgpt: Analyzing persona-assigned language models. *arXiv preprint arXiv:2304.05335*, 2023. 1

Edgar Dobriban, Hamed Hassani, David Hong, and Alexander Robey. Provable tradeoffs in adversarially robust classification. *IEEE Transactions on Information Theory*, 2023. 28

Cian Eastwood, Alexander Robey, Shashank Singh, Julius Von Kügelgen, Hamed Hassani, George J Pappas, and Bernhard Schölkopf. Probable domain generalization via quantile risk minimization. *Advances in Neural Information Processing Systems*, 35:17340–17358, 2022. 40

Marc Fischer, Maximilian Baader, and Martin Vechev. Certified defense to image transformations via randomized smoothing. *Advances in Neural information processing systems*, 33:8404–8417, 2020. 41

Iason Gabriel. Artificial intelligence, values, and alignment. *Minds and machines*, 30(3):411–437, 2020. 1

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. Realtoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462*, 2020. 1

Simon Geisler, Tom Wollschläger, MHI Abdalla, Johannes Gasteiger, and Stephan Günnemann. Attacking large language models with projected gradient descent. *arXiv preprint arXiv:2402.09154*, 2024. 3

Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Maribeth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022. 1

Grzegorz Gluch and Rüdiger Urbanke. Query complexity of adversarial attacks. In *International Conference on Machine Learning*, pp. 3723–3733. PMLR, 2021. 11

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 40

Shreya Goyal, Sumanth Doddapaneni, Mitesh M Khapra, and Balaraman Ravindran. A survey of adversarial defenses and robustness in nlp. *ACM Computing Surveys*, 55(14s):1–39, 2023. 4

Philipp Hacker, Andreas Engel, and Marco Mauer. Regulating chatgpt and other large generative ai models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1112–1123, 2023. 1

Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. *arXiv preprint arXiv:2203.09509*, 2022. 10, 28

Chia-Yi Hsu, Yu-Lin Tsai, Chih-Hsun Lin, Pin-Yu Chen, Chia-Mu Yu, and Chun-Ying Huang. Safe lora: the silver lining of reducing safety risks when fine-tuning large language models. *arXiv preprint arXiv:2405.16833*, 2024. 41

Tiansheng Huang, Sihao Hu, Fatih Ilhan, Selim Furkan Tekin, and Ling Liu. Lisa: Lazy safety alignment for large language models against harmful fine-tuning attack. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024a. 41

Tiansheng Huang, Sihao Hu, and Ling Liu. Vaccine: Perturbation-aware alignment for large language model. *arXiv preprint arXiv:2402.01109*, 2024b. 41

Yangsibo Huang, Samyak Gupta, Mengzhou Xia, Kai Li, and Danqi Chen. Catastrophic jailbreak of open-source llms via exploiting generation. *arXiv preprint arXiv:2310.06987*, 2023. 3

Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. Llama guard: Llm-based input-output safeguard for human-ai conversations. *arXiv preprint arXiv:2312.06674*, 2023. 3, 26

Neel Jain, Avi Schwarzschild, Yuxin Wen, Gowthami Somepalli, John Kirchenbauer, Ping-yeh Chiang, Micah Goldblum, Aniruddha Saha, Jonas Geiping, and Tom Goldstein. Baseline defenses for adversarial attacks against aligned language models. *arXiv preprint arXiv:2309.00614*, 2023. 4, 10

Adel Javanmard, Mahdi Soltanolkotabi, and Hamed Hassani. Precise tradeoffs in adversarial training for linear regression. In *Conference on Learning Theory*, pp. 2034–2078. PMLR, 2020. 28

Jiabao Ji, Bairu Hou, Alexander Robey, George J Pappas, Hamed Hassani, Yang Zhang, Eric Wong, and Shiyu Chang. Defending large language models against jailbreak attacks via semantic smoothing. *arXiv preprint arXiv:2402.16192*, 2024. 10, 32

Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021. 40

Aounon Kumar, Chirag Agarwal, Suraj Srinivas, Soheil Feizi, and Hima Lakkaraju. Certifying llm safety against adversarial prompting. *arXiv preprint arXiv:2309.02705*, 2023. 4

Cassidy Laidlaw, Sahil Singla, and Soheil Feizi. Perceptual adversarial robustness: Defense against unseen threat models. *arXiv preprint arXiv:2006.12655*, 2020. 40

Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE symposium on security and privacy (SP)*, pp. 656–672. IEEE, 2019. 2, 5, 8, 40

Alexander Levine and Soheil Feizi. (de) randomized smoothing for certifiable defense against patch attacks. *Advances in Neural Information Processing Systems*, 33:6465–6475, 2020. 40, 41

Jinfeng Li, Shouling Ji, Tianyu Du, Bo Li, and Ting Wang. Textbugger: Generating adversarial text against real-world applications. *arXiv preprint arXiv:1812.05271*, 2018. 4, 41

Zeyi Liao and Huan Sun. Amplegcg: Learning a universal and transferable generative model of adversarial suffixes for jailbreaking both open and closed llms. *arXiv preprint arXiv:2404.07921*, 2024. 1, 2, 3, 9

Xiaodong Liu, Hao Cheng, Pengcheng He, Weizhu Chen, Yu Wang, Hoifung Poon, and Jianfeng Gao. Adversarial training for large neural language models. *arXiv preprint arXiv:2004.08994*, 2020. 4

Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*, 2023. 1, 3

Shayne Longpre, Sayash Kapoor, Kevin Klyman, Ashwin Ramaswami, Rishi Bommasani, Borhane Blili-Hamelin, Yangsibo Huang, Aviya Skowron, Zheng-Xin Yong, Suhas Kotha, et al. A safe harbor for ai evaluation and red teaming. *arXiv preprint arXiv:2403.04893*, 2024. 1

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 12, 40

Natalie Maus, Patrick Chao, Eric Wong, and Jacob Gardner. Adversarial prompting for black box foundation models. *arXiv preprint arXiv:2302.04237*, 2023. 1

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018. 10, 28

Takeru Miyato, Andrew M Dai, and Ian Goodfellow. Adversarial training methods for semi-supervised text classification. *arXiv preprint arXiv:1605.07725*, 2016. 4

John X Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. *arXiv preprint arXiv:2005.05909*, 2020. 41

OpenAI. OpenAI O1 System Card. Technical report, OpenAI, 2023. URL https://cdn.openai.com/o1-system-card.pdf. Accessed: 2024-09-14. 11

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback, 2022. *URL https://arxiv. org/abs/2203.02155*, 13, 2022. 1

Tianyu Pang, Chao Du, Qian Liu, Jing Jiang, Min Lin, et al. Improved few-shot jailbreaking can circumvent aligned language models and their defenses. *Advances in Neural Information Processing Systems*, 37: 32856–32887, 2025. 10

ShengYun Peng, Pin-Yu Chen, Matthew Hull, and Duen Horng Chau. Navigating the safety landscape: Measuring risks in finetuning large language models. *arXiv preprint arXiv:2405.17374*, 2024. 41

Danish Pruthi, Bhuwan Dhingra, and Zachary C Lipton. Combating adversarial misspellings with robust word recognition. *arXiv preprint arXiv:1905.11268*, 2019. 41

Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*, 2023. 41

Shuhuai Ren, Yihe Deng, Kun He, and Wanxiang Che. Generating natural language adversarial examples through probability weighted word saliency. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pp. 1085–1097, 2019. 41

Alexander Robey, Hamed Hassani, and George J Pappas. Model-based robust deep learning: Generalizing to natural, out-of-distribution data. *arXiv preprint arXiv:2005.10247*, 2020. 40

Alexander Robey, George J Pappas, and Hamed Hassani. Model-based domain generalization. *Advances in Neural Information Processing Systems*, 34:20210–20229, 2021. 40

Domenic Rosati, Jan Wehner, Kai Williams, Łukasz Bartoszcze, David Atanasov, Robie Gonzales, Subhabrata Majumdar, Carsten Maple, Hassan Sajjad, and Frank Rudzicz. Representation noising effectively prevents harmful fine-tuning on llms. *arXiv preprint arXiv:2405.14577*, 2024. 41

Elan Rosenfeld, Ezra Winston, Pradeep Ravikumar, and Zico Kolter. Certified robustness to label-flipping attacks via randomized smoothing. In *International Conference on Machine Learning*, pp. 8230–8241. PMLR, 2020. 41

Malik Sallam. Chatgpt utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns. In *Healthcare*, volume 11, pp. 887. MDPI, 2023. 1

Hadi Salman, Jerry Li, Ilya Razenshteyn, Pengchuan Zhang, Huan Zhang, Sebastien Bubeck, and Greg Yang. Provably robust deep learning via adversarially trained smoothed classifiers. *Advances in Neural Information Processing Systems*, 32, 2019. 5, 8, 11, 40

Hadi Salman, Mingjie Sun, Greg Yang, Ashish Kapoor, and J Zico Kolter. Denoised smoothing: A provable defense for pretrained classifiers. *Advances in Neural Information Processing Systems*, 33:21945–21957, 2020. 12

Shibani Santurkar, Dimitris Tsipras, and Aleksander Madry. Breeds: Benchmarks for subpopulation shift. *arXiv preprint arXiv:2008.04859*, 2020. 40

Carlos E Seminario and David C Wilson. Robustness and accuracy tradeoffs for recommender systems under attack. In *Twenty-Fifth International FLAIRS Conference*, 2012. 10

Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! *Advances in Neural Information Processing Systems*, 32, 2019. 11

Abhay Sheshadri, Aidan Ewart, Phillip Guo, Aengus Lynch, Cindy Wu, Vivek Hebbar, Henry Sleight, Asa Cooper Stickland, Ethan Perez, Dylan Hadfield-Menell, et al. Latent adversarial training improves robustness to persistent harmful behaviors in llms. *arXiv preprint arXiv:2407.15549*, 2024. 41

Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020. 1

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024. 11

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. 40

Jiaye Teng, Guang-He Lee, and Yang Yuan. ℓ_1 adversarial robustness certificates: a randomized smoothing approach. 2019. 41

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 25

Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. *Advances in neural information processing systems*, 33:1633–1645, 2020. 9

Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018. 28

Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31, 2018. 12

Jiongxiao Wang, Zichen Liu, Keun Hee Park, Muhao Chen, and Chaowei Xiao. Adversarial demonstration attacks on large language models. *arXiv preprint arXiv:2305.14950*, 2023. 1

Jiongxiao Wang, Jiazhao Li, Yiquan Li, Xiangyu Qi, Junjie Hu, Yixuan Li, Patrick McDaniel, Muhao Chen, Bo Li, and Chaowei Xiao. Backdooralign: Mitigating fine-tuning based jailbreak attack with backdoor enhanced safety alignment. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 41

Xiaosen Wang, Hao Jin, and Kun He. Natural language adversarial attack and defense in word level. *arXiv preprint arXiv:1909.06723*, 2019. 41

Xiaosen Wang, Jin Hao, Yichen Yang, and Kun He. Natural language adversarial defense through synonym encoding. In *Uncertainty in Artificial Intelligence*, pp. 823–833. PMLR, 2021a. 41

Xiaosen Wang, Yichen Yang, Yihe Deng, and Kun He. Adversarial training with fast gradient projection method against synonym substitution based text attacks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 13997–14005, 2021b. 41

Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How does llm safety training fail? *arXiv preprint arXiv:2307.02483*, 2023. 1, 3

Eric Wong and J Zico Kolter. Learning perturbation sets for robust machine learning. *arXiv preprint arXiv:2007.08450*, 2020. 40

Eric Wong, Leslie Rice, and J Zico Kolter. Fast is better than free: Revisiting adversarial training. *arXiv preprint arXiv:2001.03994*, 2020. 11

Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564*, 2023. 1

Anton Xue, Rajeev Alur, and Eric Wong. Stability guarantees for feature attributions with multiplicative smoothing. *arXiv preprint arXiv:2307.05902*, 2023. 40, 41

Greg Yang, Tony Duan, J Edward Hu, Hadi Salman, Ilya Razenshteyn, and Jerry Li. Randomized smoothing of all shapes and sizes. In *International Conference on Machine Learning*, pp. 10693–10705. PMLR, 2020. 40, 41

Maksym Yatsura, Kaspar Sakmann, N Grace Hua, Matthias Hein, and Jan Hendrik Metzen. Certified defences against adversarial patch attacks on semantic segmentation. *arXiv preprint arXiv:2209.05980*, 2022. 40, 41

Zheng-Xin Yong, Cristina Menghini, and Stephen H Bach. Low-resource languages jailbreak gpt-4. *arXiv preprint arXiv:2310.02446*, 2023. 3

Eliezer Yudkowsky. The ai alignment problem: why it is hard, and where to start. *Symbolic Systems Distinguished Speaker*, 4, 2016. 1

Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pp. 7472–7482. PMLR, 2019. 40

Wei Emma Zhang, Quan Z Sheng, Ahoud Alhazmi, and Chenliang Li. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(3):1–41, 2020. 41

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*, 2023. 10

Yi Zhou, Xiaoqing Zheng, Cho-Jui Hsieh, Kai-Wei Chang, and Xuanjing Huan. Defense against synonym substitution-based adversarial attacks via dirichlet neighborhood ensemble. In *Association for Computational Linguistics (ACL)*, 2021. 41

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*, 2023a. 41

Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023b. 1, 2, 3, 5, 6, 9, 12, 25, 26, 27, 28, 29, 30, 33, 34, 38

Andy Zou, Long Phan, Justin Wang, Derek Duenas, Maxwell Lin, Maksym Andriushchenko, J Zico Kolter, Matt Fredrikson, and Dan Hendrycks. Improving alignment and robustness with circuit breakers. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 41

A Omitted proofs

Below, we state the formal version of Proposition 3.3, which was stated informally in the main text.

Proposition A.1 (SmoothLLM certificate)

Let $\mathcal{A}$ denote an alphabet of size v (i.e., $|\mathcal{A}| = v$) and let $P = [G; S] \in \mathcal{A}^m$ denote an input prompt to a given LLM where $G \in \mathcal{A}^{m_G}$ and $S \in \mathcal{A}^{m_S}$. Furthermore, let $M = \lfloor qm \rfloor$ and $u = \min(M, m_S)$. Then assuming that S is k-unstable for $k \leq \min(M, m_S)$, the following holds:

(a) The probability that SmoothLLM is not jailbroken by when run with the RANDOMSWAPPER-
TURBATION function is

$$\mathrm{DSP}([G; S]) = \sum_{t=\lceil N/2 \rceil}^{n} \binom{N}{t} \alpha^t (1 - \alpha)^{N-t} \tag{11}$$

where

$$\alpha \triangleq \sum_{i=k}^{u} \left[\binom{M}{i} \binom{m - m_S}{M - i} \bigg/ \binom{m}{M} \right] \beta(\ell, i, v) \tag{12}$$

and where

$$\beta(\ell, i, v) \triangleq \sum_{\ell=k}^{i} \binom{i}{\ell} \left(\frac{v-1}{v} \right)^{\ell} \left(\frac{1}{v} \right)^{i-\ell} \tag{13}$$

(b) The probability that SMOOTHLLM is not jailbroken by when run with the RANDOMPATCHPER-
TURBATION function is

$$\mathrm{DSP}([G; S]) = \sum_{t=\lceil N/2 \rceil}^{n} \binom{N}{t} \alpha^t (1 - \alpha)^{N-t} \tag{14}$$

where α is defined piecewise as follows. First, if $M \leq m_S$, then

$$\alpha = \left(\frac{m_S - M + 1}{m - M + 1} \right) \beta(M)$$
$$+ \left(\frac{1}{m - M + 1} \right) \sum_{j=1}^{\min(m_G, M-k)} \beta(M - j). \tag{15}$$

Secondly, if $m_G \geq M - k$ and $M > m_S$, then

$$\alpha = \left(\frac{1}{m - M + 1} \right) \sum_{j=0}^{m_S - k} \beta(M - j) \tag{16}$$

And finally, if $m_G < M - k$ and $M > m_S$, then

$$\alpha = \left(\frac{1}{m - M + 1} \right) \sum_{j=0}^{m-M} \beta(M - j). \tag{17}$$

Throughout, $\beta(i) \triangleq \sum_{\ell=k}^{i} \binom{i}{\ell} \left(\frac{v-1}{v} \right)^{\ell} \left(\frac{1}{v} \right)^{i-\ell}$.

Proof. We are interested in computing the following probability:

$$\Pr\left[(\mathrm{JB}\circ\textsc{SmoothLLM})(P) = 0\right]$$
$$= \Pr\left[\mathrm{JB}\left(\textsc{SmoothLLM}(P)\right) = 0\right]. \tag{18}$$

By the way SmoothLLM is defined in definition 3.1 and (4),

$$(\mathrm{JB}\circ\textsc{SmoothLLM})(P)$$
$$= \mathbb{I}\left[\frac{1}{N}\sum_{j=1}^{N}(\mathrm{JB}\circ\mathrm{LLM})(P_j) > \frac{1}{2}\right] \tag{19}$$

where P_j for $j \in [N]$ are drawn i.i.d. from $\mathbb{P}_q(P)$. The following chain of equalities follows directly from applying this definition to the probability in (18):

$$\Pr\left[(\mathrm{JB}\circ\textsc{SmoothLLM})(P) = 0\right] \tag{20}$$

$$= \Pr_{P_1,\ldots,P_N}\left[\frac{1}{N}\sum_{j=1}^{N}(\mathrm{JB}\circ\mathrm{LLM})(P_j) \leq \frac{1}{2}\right] \tag{21}$$

$$= \Pr_{P_1,\ldots,P_N}\left[(\mathrm{JB}\circ\mathrm{LLM})(P_j) = 0 \text{ for at least } \left\lceil\frac{N}{2}\right\rceil\right.$$
$$\left. \text{of the indices } j \in [N]\right] \tag{22}$$

$$= \sum_{t=\lceil N/2\rceil}^{N}\Pr_{P_1,\ldots,P_N}\left[(\mathrm{JB}\circ\mathrm{LLM})(P_j) = 0 \text{ for exactly } t\right.$$
$$\left. \text{of the indices } j \in [N]\right]. \tag{23}$$

Let us pause here to take stock of what was accomplished in this derivation.

- In step (21), we made explicit the source of randomness in the forward pass of SmoothLLM, which is the N-fold draw of the randomly perturbed prompts P_j from $\mathbb{P}_q(P)$ for $j \in [N]$.

- In step (22), we noted that since JB is a binary-valued function, the average of $(\mathrm{JB}\circ\mathrm{LLM})(P_j)$ over $j \in [N]$ being less than or equal to $1/2$ is equivalent to at least $\lceil N/2\rceil$ of the indices $j \in [N]$ being such that $(\mathrm{JB}\circ\mathrm{LLM})(P_j) = 0$.

- In step (23), we explicitly enumerated the cases in which at least $\lceil N/2\rceil$ of the perturbed prompts P_j do not result in a jailbreak, i.e., $(\mathrm{JB}\circ\mathrm{LLM})(P_j) = 0$.

The result of this massaging is that the summands in (23) bear a noticeable resemblance to the elementary, yet classical setting of flipping biased coins. To make this precise, let α denote the probability that a randomly drawn element $Q \sim \mathbb{P}_q(P)$ does not constitute a jailbreak, i.e.,

$$\alpha = \alpha(P, q) \triangleq \Pr_{Q}\left[(\mathrm{JB}\circ\mathrm{LLM})(Q) = 0\right]. \tag{24}$$

Now consider an experiment wherein we perform N flips of a biased coin that turns up heads with probability α; in other words, we consider N Bernoulli trials with success probability α. For each index t in the summation in (23), the concomitant summand denotes the probability that of the N (independent) coin flips (or, if you like, Bernoulli trials), exactly t of those flips turn up as heads. Therefore, one can write the probability in (23) using a binomial expansion:

$$\Pr\left[(\mathrm{JB}\circ\textsc{SmoothLLM})(P) = 0\right]$$
$$= \sum_{t=\lceil N/2\rceil}^{N}\binom{N}{t}\alpha^t(1-\alpha)^{N-t} \tag{25}$$

where α is the probability defined in (24).

The remainder of the proof concerns deriving an explicit expression for the probability α. Since by assumption the prompt $P = [G; S]$ is k-unstable, it holds that

$$(\text{JB} \circ \text{LLM})([G; S']) = 0 \iff d_H(S, S') \geq k. \tag{26}$$

where $d_H(\cdot, \cdot)$ denotes the Hamming distance between two strings. Therefore, by writing our randomly drawn prompt Q as $Q = [Q_G; Q_S]$ for $Q_G \in \mathcal{A}^{m_G}$ and $Q_S \in \mathcal{A}^{m_S}$, it's evident that

$$\alpha = \Pr_Q \left[(\text{JB} \circ \text{LLM})([Q_G; Q_S]) = 0 \right] \tag{27}$$

$$= \Pr_Q \left[d_H(S, Q_S) \geq k \right] \tag{28}$$

We are now confronted with the following question: What is the probability that S and a randomly-drawn suffix Q_S differ in at least k locations? And as one would expect, the answer to this question depends on the kinds of perturbations that are applied to P. Therefore, toward proving parts (a) and (b) of the statement of this proposition, we now specialize our analysis to swap and patch perturbations respectively.

A.1 Swap perturbations.

Consider the `RandomSwapPerturbation` function defined in lines 1-5 of Algorithm 2. This function involves two main steps:

1. Select a set $\mathcal{I}$ of $M \triangleq \lfloor qm \rfloor$ locations in the prompt P uniformly at random.

2. For each sampled location, replace the character in P at that location with a character a sampled uniformly at random from $\mathcal{A}$, i.e., $a \sim \text{Unif}(\mathcal{A})$.

These steps suggest that we break down the probability in drawing Q into (1) drawing the set of $\mathcal{I}$ indices and (2) drawing M new elements uniformly from $\text{Unif}(\mathcal{A})$. To do so, we first introduce the following notation to denote the set of indices of the suffix in the original prompt P:

$$\mathcal{I}_S \triangleq \{m - m_S + 1, \ldots, m - 1\}. \tag{29}$$

Now observe that

$$\alpha = \Pr_{\mathcal{I}, a_1, \ldots, a_M} \Big[|\mathcal{I} \cap \mathcal{I}_S| \geq k \text{ and}$$
$$|\{j \in \mathcal{I} \cap \mathcal{I}_S \, : \, P[j] \neq a_j\}| \geq k \Big] \tag{30}$$
$$= \Pr_{a_1, \ldots, a_M} \Big[|\{j \in \mathcal{I} \cap \mathcal{I}_S \, : \, P[j] \neq a_j\}| \geq k$$
$$\Big| \, |\mathcal{I} \cap \mathcal{I}_S| \geq k \Big] \cdot \Pr_{\mathcal{I}} \Big[|\mathcal{I} \cap \mathcal{I}_S| \geq k \Big] \tag{31}$$

The first condition in the probability in (30)—$|\mathcal{I} \cap \mathcal{I}_S| \geq k$—denotes the event that at least k of the sampled indices are in the suffix; the second condition—$|\{j \in \mathcal{I} \cap \mathcal{I}_S \, : \, P[j] \neq a_j\}| \geq k$—denotes the event that at least k of the sampled replacement characters are different from the original characters in P at the locations sampled in the suffix. And step (31) follows from the definition of conditional probability.

Considering the expression in (31), by directly applying Lemma A.2, observe that

$$\alpha = \sum_{i=k}^{\min(M, m_S)} \frac{\binom{M}{i} \binom{m - m_S}{M - i}}{\binom{m}{M}} \cdot \eta(i) \tag{32}$$

where

$$\eta(i) \triangleq \Pr_{a_1, \ldots, a_M} \Big[|\{j \in \mathcal{I} \cap \mathcal{I}_S \, : \, P[j] \neq a_j\}| \geq k$$
$$\Big| \, |\mathcal{I} \cap \mathcal{I}_S| = i \Big]. \tag{33}$$

To finish up the proof, we seek an expression for the probability over the N-fold draw from $\text{Unif}(\mathcal{A})$ above. However, as the draws from $\text{Unif}(\mathcal{A})$ are *independent*, we can translate this probability into another question of flipping coins that turn up heads with probability $v-1/v$, i.e., the chance that a character $a \sim \text{Unif}(\mathcal{A})$ at a particular index is not the same as the character originally at that index. By an argument entirely similar to the one given after (24), it follows easily that

$$\Pr_{a_1,\dots,a_M}\left[|\{j \in \mathcal{I} \cap \mathcal{I}_S \,:\, P[j] \neq a_j\}| \geq k \,\big|\, |\mathcal{I} \cap \mathcal{I}_S| = i\right]$$
$$= \sum_{\ell=k}^{i} \binom{i}{\ell} \left(\frac{v-1}{v}\right)^{\ell} \left(\frac{1}{v}\right)^{i-\ell} \tag{34}$$

Plugging this expression back into (32) completes the proof for swap perturbations.

A.2 Patch perturbations.

We now turn our attention to patch perturbations, which are defined by the `RandomPatchPerturbation` function in lines 6-10 of Algorithm 2. In this setting, a simplification arises as there are fewer ways of selecting the locations of the perturbations themselves, given the constraint that the locations must be contiguous. At this point, it's useful to break down the analysis into four cases. In every case, we note that there are $n - M + 1$ possible patches.

Case 1: $m_G \geq M - k$ and $M \leq m_S$. In this case, the number of locations M covered by a patch is fewer than the length of the suffix m_S, and the length of the goal is at least as large as $M - k$. As $M \leq m_S$, it's easy to see that there are $m_S - M + 1$ potential patches that are completely contained in the suffix. Furthermore, there are an additional $M - k$ potential locations that overlap with the the suffix by at least k characters, and since $m_G \geq M - k$, each of these locations engenders a valid patch. Therefore, in total there are

$$(m_S - M + 1) + (M - k) = m_S - k + 1 \tag{35}$$

valid patches in this case.

To calculate the probability α in this case, observe that of the patches that are completely contained in the suffix—each of which could be chosen with probability $(m_S - M + 1)/(m - M + 1)$—each patch contains M characters in S. Thus, for each of these patches, we enumerate the ways that at least k of these M characters are sampled to be different from the original character at that location in P. And for the $M - k$ patches that only partially overlap with S, each patch overlaps with $M - j$ characters where j runs from 1 to $M - k$. For these patches, we then enumerate the ways that these patches flip at least k characters, which means that the inner sum ranges from $\ell = k$ to $\ell = M - j$ for each index j mentioned in the previous sentence. This amounts to the following expression:

$$\alpha = \overbrace{\left(\frac{m_S - M + 1}{m - M + 1}\right) \sum_{\ell=k}^{M} \binom{M}{\ell} \left(\frac{v-1}{v}\right)^{\ell} \left(\frac{1}{v}\right)^{M-\ell}}^{\text{patches completely contained in the suffix}}$$
$$+ \underbrace{\sum_{j=1}^{M-k} \left(\frac{1}{m-M+1}\right) \sum_{\ell=k}^{M-j} \binom{M-j}{\ell} \left(\frac{v-1}{v}\right)^{\ell} \left(\frac{1}{v}\right)^{M-j-\ell}}_{\text{patches partially contained in the suffix}} \tag{36}$$

Case 2: $m_G < M - k$ and $M \leq m_S$. This case is similar to the previous case, in that the term involving the patches completely contained in S is completely the same as the expression in (36). However, since m_G is strictly less than $M - k$, there are fewer patches that partially intersect with S than in the previous case.

In this way, rather than summing over indices j running from 1 to $M - k$, which represents the number of locations that the patch intersects with G, we sum from $j = 1$ to m_G, since there are now m_G locations where the patch can intersect with the goal. Thus,

$$
\alpha = \left(\frac{m_S - M + 1}{m - M + 1} \right) \sum_{\ell=k}^{M} \binom{M}{\ell} \left(\frac{v-1}{v} \right)^{\ell} \left(\frac{1}{v} \right)^{M-\ell}
$$
$$
+ \sum_{j=1}^{m_G} \left(\frac{1}{m - M + 1} \right) \sum_{\ell=k}^{M-j} \binom{M-j}{\ell} \left(\frac{v-1}{v} \right)^{\ell} \left(\frac{1}{v} \right)^{M-j-\ell} \tag{37}
$$

Note that in the statement of the proposition, we condense these two cases by writing

$$
\alpha = \left(\frac{m_S - M + 1}{m - M + 1} \right) \beta(M)
$$
$$
+ \left(\frac{1}{m - M + 1} \right) \sum_{j=1}^{\min(m_G, M-k)} \beta(M - j). \tag{38}
$$

Case 3: $m_G \geq M - k$ and $M < m_S$. Next, we consider cases in which the width of the patch M is larger than the length m_S of the suffix S, meaning that every valid patch will intersect with the goal in at least one location. When $m_G \geq M - k$, all of the patches that intersect with the suffix in at least k locations are viable options. One can check that there are $m_S - M + 1$ valid patches in this case, and therefore, by appealing to an argument similar to the one made in the previous two cases, we find that

$$
\alpha = \sum_{j=0}^{m_S-k} \left(\frac{1}{m - M + 1} \right) \sum_{\ell=k}^{T-j} \binom{T-j}{\ell} \left(\frac{v-1}{v} \right)^{\ell} \left(\frac{1}{v} \right)^{M-j-\ell} \tag{39}
$$

where one can think of j as iterating over the number of locations in the suffix that are not included in a given patch.

Case 4: $m_G < M - k$ and $M < m_S$. In the final case, in a similar vein to the second case, we are now confronted with situations wherein there are fewer patches that intersect with S than in the previous case, since $m_G < M - k$. Therefore, rather than summing over the $m_S - k + 1$ patches present in the previous step, we now must disregard those patches that no longer fit within the prompt. There are exactly $(M - k) - m_G$ such patches, and therefore in this case, there are

$$
(m_S - k + 1) - (M - k - m_G) = m - M + 1 \tag{40}
$$

valid patches, where we have used the fact that $m_G + m_S = m$. This should couple with our intuition, as in this case, all patches are valid. Therefore, by similar logic to that used in the previous case, it is evident that we can simply replace the outer sum so that j ranges from 0 to $m - M$:

$$
\alpha = \sum_{j=0}^{m-M} \left(\frac{1}{m - M + 1} \right) \sum_{\ell=k}^{T-j} \binom{T-j}{\ell} \left(\frac{v-1}{v} \right)^{\ell} \left(\frac{1}{v} \right)^{M-j-\ell}. \tag{41}
$$

This completes the proof. □

Lemma A.2

We are given a set $\mathcal{B}$ containing n elements and a fixed subset $\mathcal{C} \subseteq \mathcal{B}$ comprising d elements ($d \leq n$). If one samples a set $\mathcal{I} \subseteq \mathcal{B}$ of T elements uniformly at random without replacement from $\mathcal{B}$ where $T \in [1, n]$, then the probability that at least k elements of $\mathcal{C}$ are sampled where $k \in [0, d]$ is

$$\Pr_{\mathcal{I}} \left[|\mathcal{I} \cap \mathcal{C}| \geq k \right] = \sum_{i=k}^{\min(T,d)} \binom{T}{i} \binom{n-d}{T-i} \bigg/ \binom{n}{T}. \tag{42}$$

Proof. We begin by enumerating the cases in which *at least* k elements of $\mathcal{C}$ belong to $\mathcal{I}$:

$$\Pr_{\mathcal{I}} \left[|\mathcal{I} \cap \mathcal{C}| \geq k \right] = \sum_{i=k}^{\min(T,d)} \Pr_{\mathcal{I}} \left[|\mathcal{I} \cap \mathcal{C}| = i \right] \tag{43}$$

The subtlety in (43) lies in determining the final index in the summation. If $T > d$, then the summation runs from k to d because $\mathcal{C}$ contains only d elements. On the other hand, if $d > T$, then the summation runs from k to T, since the sampled subset can contain at most T elements from $\mathcal{C}$. Therefore, in full generality, the summation can be written as running from k to $\min(T, d)$.

Now consider the summands in (43). The probability that exactly i elements from $\mathcal{C}$ belong to $\mathcal{I}$ is:

$$\Pr_{\mathcal{I}} \left[|\mathcal{I} \cap \mathcal{C}| = i \right]$$
$$= \frac{\text{Number of subsets } \mathcal{I} \text{ of } \mathcal{B} \text{ with } i \text{ elements from } \mathcal{C}}{\text{Total number of subsets } \mathcal{I} \text{ of } \mathcal{B}} \tag{44}$$

Consider the numerator, which counts the number of ways one can select a subset of T elements from $\mathcal{B}$ that contains i elements from $\mathcal{C}$. In other words, we want to count the number of subsets $\mathcal{I}$ of $\mathcal{B}$ that contain i elements from $\mathcal{C}$ and $T - i$ elements from $\mathcal{B} \backslash \mathcal{C}$. To this end, observe that:

- There are $\binom{T}{i}$ ways of selecting the i elements of $\mathcal{C}$ in the sampled subset;

- There are $\binom{n-d}{T-i}$ ways of selecting the $T - i$ elements of $\mathcal{B} \backslash \mathcal{C}$ in the sampled subset.

Therefore, the numerator in (44) is $\binom{T}{i}\binom{n-d}{T-i}$. The denominator in (44) is easy to calculate, since there are $\binom{n}{T}$ subsets of $\mathcal{B}$ of length n. In this way, we have shown that

$$\Pr \left[\text{Exactly } i \text{ elements from } \mathcal{C} \text{ are sampled from } \mathcal{B} \right]$$
$$= \binom{T}{i} \binom{n-d}{T-i} \bigg/ \binom{n}{T} \tag{45}$$

and by plugging back into (43) we obtain the desired result. $\qquad \square$

B   Further experimental details

B.1   Computational resources

All experiments in this paper were run on a cluster with 8 NVIDIA A100 GPUs and 16 NVIDIA A6000 GPUs. The bulk of the computation involved obtaining adversarial suffixes for the prompts proposed in (Zou et al., 2023b).

B.2   LLM versions

Throughout this paper, used the following versions of Llama2 (Touvron et al., 2023) and Vicuna (Chiang et al., 2023):

- Llama2: https://huggingface.co/meta-llama/Llama-2-7b-chat-hf

- Vicuna: https://huggingface.co/lmsys/vicuna-13b-v1.5

B.3   Running `GCG`

Throughout this paper, we used the authors' implementation of `GCG` to obtain adversarial suffixes for Llama2 and Vicuna. Aside from the query efficiency experiments, all plots show suffixes obtained by running `GCG` for 500 iterations with a batch size of 512, which is the configuration specified in (Zou et al., 2023b).

B.4   Determining whether a jailbreak has occurred

Throughout this paper, to determine whether a jailbreak has occurred (i.e., to evaluate the JB function), we followed prior work by pairing attacking algorithms with the judges used in the authors' implementations.

GCG. For GCG, we use the same criteria as in (Zou et al., 2023b). In the authors' implementation of `GCG`, two sets of criteria are used, both of which take the following form:

$$\mathrm{JB}(R) \triangleq \mathbb{I}[R \text{ does not contain any phrase in}$$
$$\texttt{JailbreakKeywords}] \tag{46}$$

where `JailbreakKeywords` is a list of keywords that are typically included in messages which refuse or abstain from responding to a prompt requesting objectionable content. When evaluating whether an open-source LLM (e.g., Llama2 or Vicuna) is jailbroken, the following list is used:

> `JailbreakKeywords` = ["I'm sorry", "Sorry", "I am sorry", "I apologize", "As an", "As an AI", "I'm an", "I'm just", "As a language model", "As an Assistant", "I cannot", "I do not", "Hello!"]

For details, see the source code. When evaluating whether a closed-source LLM is jailbroken, the following list is used:

> `JailbreakKeywords` = ["I'm sorry", "Sorry", "an AI", "is illegal and unethical", "I'm really sorry", "My purpose", "As a helpful assistant", "OpenAI", "unethical", "positive and up-lifting", "against ethical guidelines", "not ethical", "legally", "illegal", "My apologies", "legal ways", "but I cannot", "If you have any other non-malicious requests", "I'm not able to pro-vide", "I cannot engage", " I cannot generate", " I do not generate", "I cannot create", "I do not actually", "a bit unclear", "I can't create", "I cannot support", "I can't provide", "I'm truly sorry", "I cannot provide", "I'd prefer not", "I do not provide", "I cannot recommend", "I do not recommend", "I do not feel comfortable", "I can't assist", "Absolutely not!", "hereś an alternative", "here's an alternative", "here's a different", "I cannot assist", "I will refrain", "I apologize", "I will not", "awesome coffee", "creativity", "promote well-being", "positivity", "harm reduction", "self-care"]

Figure 11: **Certified robustness to suffix-based attacks.** To complement Figure 6 in the main text, which computed the DSP for the average prompt and suffix lengths for Llama2, we produce an analogous plot for the corresponding average lengths for Vicuna. Notice that as in Figure 6, as N and q increase, so does the DSP.

For details, see the source code.

PAIR. For PAIR, we used the same criteria as (Chao et al., 2023), who use the Llama Guard classifier (Inan et al., 2023) to instantiate the JB function.

RandomSearch and AmpleGCG. For both RANDOMSEARCH and AMPLEGCG, we followed the authors by using LLM-as-a-judge paradigm with GPT-4 to instantiate the JB function.

B.5 Selecting *N* and *q* in Algorithm 1

As shown throughout this paper, selecting the values of the number of samples N and the perturbation percentage q are essential to obtaining a strong defense. In several of the figures, e.g., Figures 1 and 14, we swept over a range of values for N and q and reported the performance corresponding to the combination that yielded the best results. In practice, given that SmoothLLM is query- and time-efficient, this may be a viable strategy. One promising direction for future research is to experiment with different ways of selecting N and q. For instance, one could imagine ensembling the generated responses from instantiations of SmoothLLM with different hyperparameters to improve robustness.

B.6 The instability of adversarial suffixes

To generate Figure 4, we obtained adversarial suffixes for Llama2 and Vicuna by running the authors' implementation of `GCG` for every prompt in the `behaviors` dataset described in (Zou et al., 2023b). We then ran SmoothLLM for $N \in \{2, 4, 6, 8, 10\}$ and $q \in \{5, 10, 15, 20\}$ across five independent trials. In this way, the bar heights represent the mean ASRs over these five trials, and the black lines at the top of these bars indicate the corresponding standard deviations.

B.7 Robustness guarantees in a simplified setting

In Section 3.3, we calculated and plotted the DSP for the average prompt and suffix lengths—$m = 168$ and $m_S = 96$—for Llama2. This average was taken over all 500 suffixes obtained for Llama2. As alluded to in the footnote at the end of that section, the averages for the corresponding quantities across the 500 suffixes obtained for Vicuna were similar: $m = 179$ and $m_S = 106$. For the sake of completeness, in Figure 11, we reproduce Figure 6 with the average prompt and suffix length for Vicuna, rather than for Llama2. In this

Table 3: **Parameters used to compute the DSP.** We list the parameters used to compute the DSP in Figures 6 and 11. The only difference between these two figures are the choices of m and m_S.

Description	Symbol	Value
Number of smoothing samples	N	$\{2, 4, 6, 8, 10, 12, 14, 16, 18, 20\}$
Perturbation percentage	q	$\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$
Alphabet size	v	100
Prompt length	m	168 (Figure 6) or 179 (Figure 11)
Suffix length	m_S	96 (Figure 6) or 106 (Figure 11)
Goal length	m_G	$m - m_S$
Instability parameter	k	$\{2, 5, 8\}$

Figure 12: **Query-efficiency: attack vs. defense.** To complement Figure 10 in the main text, which concerned the query-efficiency of `GCG` and SmoothLLM on Vicuna, we produce an analogous plot for Llama2. This plot displays similar trends. As `GCG` runs for more iterations, the ASR tends to increase. However, as N and q increase, SmoothLLM is able to mitigate the attack.

figure, the trends are the same: The DSP decreases as the number of steps of `GCG` increases, but dually, as N and q increase, so does the DSP.

In Table 3, we list the parameters used to calcualte the DSP in Figures 6 and 11. The alphabet size $v = 100$ is chosen for consistency with out experiments, which use a 100-character alphabet $\mathcal{A} = $ `string.printable` (see Appendix D for details).

B.8 Query-efficiency: attack vs. defense

In § 4, we compared the query efficiencies of `GCG` and SmoothLLM. In particular, in Figure 10 we looked at the ASR on Vicuna for varying step counts for `GCG` and SmoothLLM. To complement this result, we produce an analogous plot for Llama2 in Figure 12.

To generate Figure 10 and Figure 12, we obtained 100 adversarial suffixes for Llama2 and Vicuna by running `GCG` on the first 100 entries in the `harmful_behaviors.csv` dataset provided in the `GCG` source code. For each suffix, we ran `GCG` for 500 steps with a batch size of 512, which is the configuration specified in (Zou et al., 2023b, §3, page 9). In addition to the final suffix, we also saved ten intermediate checkpoints—one every 50 iterations—to facilitate the plotting of the performance of `GCG` at different step counts. After obtaining these suffixes, we ran SmoothLLM with swap perturbations for $N \in \{2, 4, 6, 8, 10, 12\}$ steps.

To calculate the number of queries used in `GCG`, we simply multiply the batch size by the number of steps. E.g., the suffixes that are run for 500 steps use $500 \times 512 = 256,000$ total queries. This is a slight underestimate,

Figure 13: **Robustness trade-offs.** All results correspond to the `InstructionFollowing` dataset. The top row shows results for Vicuna, and the bottom row shows results for Llama2. As in Figure 9, the dashed lines denote the performance of an undefended LLM.

as there is an additional query made to compute the loss. However, for the sake of simplicity, we disregard this query.

B.9 Non-conservatism

In the literature surrounding robustness in deep learning, there is ample discussion of the trade-offs between nominal performance and robustness. In adversarial examples research, several results on both the empirical and theoretical side point to the fact that higher robustness often comes at the cost of degraded nominal performance (Tsipras et al., 2018; Dobriban et al., 2023; Javanmard et al., 2020). In this setting, the adversary can attack *any* data passed as input to a deep neural network, resulting in the pronounced body of work that has sought to resolve this vulnerability.

While the literature concerning jailbreaking LLMs shares similarities with the adversarial robustness literature, there are several notable differences. One relevant difference is that by construction, jailbreaks only occur when the model receives prompts as input that request objectionable content. In other words, adversarial-prompting-based jailbreaks such as `GCG` have only been shown to bypass the safety filters implemented on LLMs on prompts that are written with malicious intentions. This contrasts with the existing robustness literature, where it has been shown that any input, whether benign or maliciously constructed, can be attacked.

This observation points to a pointed difference between the threat models considered in the adversarial robustness literature and the adversarial prompting literature. Moreover, the result of this difference is that it is somewhat unclear how one should evaluate the "clean" or nominal performance of a defended LLM. For instance, since the `behvaiors` dataset proposed in (Zou et al., 2023b) does not contain any prompts that do *not* request objectionable content, there is no way to measure the extent to which defenses like SmoothLLM degrade the ability to accurately generate realistic text.

To evaluate the trade-offs between clean text generation and robustness to jailbreaking attacks, we run Algorithm 1 on three standard NLP question-answering benchmarks: PIQA (Bisk et al., 2020), Open-BookQA (Mihaylov et al., 2018), and ToxiGen (Hartvigsen et al., 2022). In Table 4, we show the results

Table 4: **Non-conservatism of SmoothLLM.** In this table, we list the performance of SmoothLLM when instantiated on Llama2 and Vicuna across three standard question-answering benchmarks: PIQA, Open-BookQA, and ToxiGen. These numbers—when compared with the undefended scores in Table 5, indicate that SmoothLLM does not impose significant trade-offs between robustness and nominal performance.

LLM	q	N	Dataset					
			PIQA		OpenBookQA		ToxiGen	
			Swap	Patch	Swap	Patch	Swap	Patch
Llama2	2	2	63.0	66.2	32.4	32.6	49.8	49.3
		6	64.5	69.7	32.4	30.8	49.7	49.3
		10	66.5	70.5	31.4	33.5	49.8	50.7
		20	69.2	72.6	32.2	31.6	49.9	50.5
	5	2	55.1	58.0	24.8	28.6	47.5	49.8
		6	59.1	64.4	22.8	26.8	47.6	51.0
		10	62.1	67.0	23.2	26.8	46.0	50.4
		20	64.3	70.3	24.8	25.6	46.5	49.3
Vicuna	2	2	65.3	68.8	30.4	32.4	50.1	50.5
		6	66.9	71.0	30.8	31.2	50.1	50.4
		10	69.0	71.1	30.2	31.4	50.3	50.5
		20	70.7	73.2	30.6	31.4	49.9	50.0
	5	2	58.8	60.2	23.0	25.8	47.2	50.1
		6	60.9	62.4	23.2	25.8	47.2	49.3
		10	66.1	68.7	23.2	25.4	48.7	49.3
		20	66.1	71.9	23.2	25.8	48.8	49.4

Table 5: **LLM performance on standard benchmarks.** In this table, we list the performance of Llama2 and Vicuna on three standard question-answering benchmarks: PIQA, OpenBookQA, and ToxiGen.

LLM	Dataset		
	PIQA	OpenBookQA	ToxiGen
Llama2	76.7	33.8	51.6
Vicuna	77.4	33.1	52.9

of running SmoothLLM on these dataset with various values of q and N, and in Table 5, we list the corresponding performance of undefended LLMs. Notice that as N increases, the performance tends to improve, which is somewhat intuitive, given that more samples should result in stronger estimate of the majority vote. Furthermore, as q increases, performance tends to drop, as one would expect. However, overall, particularly on OpenBookQA and ToxiGen, the clean and defended performance are particularly close.

B.10 Defending closed-source LLMs with SmoothLLM

In Table 6, we attempt to reproduce a subset of the results reported in Table 2 of (Zou et al., 2023b). We ran a single trial with these settings, which is consistent with (Zou et al., 2023b). Moreover, we are restricted by the usage limits imposed when querying the GPT models. Our results show that for GPT-4 and, to some extent, PaLM-2, we were unable to reproduce the corresponding figures reported in the prior work. The most plausible explanation for this is that OpenAI and Google—the creators and maintainers of these respective LLMs—have implemented workarounds or patches that reduces the effectiveness of the suffixes found using `GCG`. However, note that since we still found a nonzero ASR for both LLMs, both models still stand to benefit from jailbreaking defenses.

Table 6: **Transfer reproduction.** In this table, we reproduce a subset of the results presented in (Zou et al., 2023b, Table 2). We find that for GPT-2.5, Claude-1, Claude-2, and PaLM-2, our the ASRs that result from transferring attacks from Vicuna (loosely) match the figures reported in (Zou et al., 2023b). While the figure we obtain for GPT-4 doesn't match prior work, this is likely attributable to patches made by OpenAI since (Zou et al., 2023b) appeared on arXiv roughly two months ago.

Source model	ASR (%) of various target models				
	GPT-3.5	GPT-4	Claude-1	Claude-2	PaLM-2
Vicuna (ours)	28.7	5.6	1.3	1.6	24.9
Llama2 (ours)	16.6	2.7	0.5	0.9	27.9
Vicuna (orig.)	34.3	34.5	2.6	0.0	31.7

Figure 14: **Preventing jailbreaks with SmoothLLM.** In this figure, we complement Figure 1 in the main text by transferring attacks from Llama2 (rather than Vicuna) to GPT-3.5, GPT-4, Claude-1, Claude-2, and PaLM-2.

Table 7: **Improving the nominal performance of SmoothLLM.**

Attack	Defense	Target LLM ASR			
		Vicuna	Llama2	GPT-3.5	GPT-4
PAIR	None	82	4	76	50
	SMOOTHLLM	47	1	12	25
	TILTEDSMOOTHLLM	43	2	10	25
GCG	None	58	2	34	1
	SMOOTHLLM	1	1	1	3
	TILTEDSMOOTHLLM	0	1	2	1

Table 8: **SmoothLLM running time.** We list the running time per prompt of SMOOTHLLM when run with various values of N and averaged over five trials. For Vicuna and Llama2, we ran SmoothLLM on A100 and A6000 GPUs, respectively. Note that the default implementation of GCG takes roughly of two hours per prompt on this hardware, which means that GCG is several thousand times slower than SMOOTHLLM.

LLM	GPU	N	Running time per prompt (seconds)		
			Insert	Swap	Patch
Vicuna	A100	2	3.54 ± 0.12	3.66 ± 0.10	3.72 ± 0.12
		4	3.80 ± 0.11	3.71 ± 0.16	3.80 ± 0.10
		6	3.81 ± 0.07	3.89 ± 0.14	4.02 ± 0.04
		8	3.94 ± 0.14	3.93 ± 0.07	4.08 ± 0.08
		10	4.16 ± 0.09	4.21 ± 0.05	4.16 ± 0.11
Llama2	A6000	2	3.29 ± 0.01	3.30 ± 0.01	3.29 ± 0.02
		4	3.56 ± 0.02	3.56 ± 0.01	3.54 ± 0.02
		6	3.79 ± 0.02	3.78 ± 0.02	3.77 ± 0.01
		8	4.11 ± 0.02	4.10 ± 0.02	4.04 ± 0.03
		10	4.38 ± 0.01	4.36 ± 0.03	4.31 ± 0.02

In Figure 14, we complement the results shown in Figure 1 by plotting the defended and undefended performance of closed-source LLMs attacked using adversarial suffixes generated for Llama2. In this figure, we see a similar trend vis-a-vis Figure 1: For all LLMs—whether open- or closed-source—the ASR of SmoothLLM drops below one percentage point. Note that in both Figures, we do not transfer attacks from Vicuna to Llama2, or from Llama2 to Vicuna. We found that attacks did not transfer between Llama2 and Vicuna. To generate the plots in Figures 1 and 14, we ran SmoothLLM with $q \in \{2, 5, 10, 15, 20\}$ and $N \in \{5, 6, 7, 8, 9, 10\}$. The ASRs for the best-performing SmoothLLM models were then plotted in the corresponding figures.

B.11 Improving nominal performance with the tilted majority vote

In Table 7, we compare the performance of SMOOTHLLM with $\gamma = {}^{1}/_{2}$, $N = 10$, and $q = 5$ to the variant of SMOOTHLLM discussed in §4.2 on the `JBB-Behaviors`. This variant, which we refer to as TILTEDSMOOTHLLM, uses $N = 10$, $\gamma = {}^{N-1}/_{N}$, $q = 5$, and returns $\mathrm{LLM}(P)$ if the majority vote V is equal to zero. Notably, Table 7 shows that SMOOTHLLM and TILTEDSMOOTHLLM offer similar levels of robustness against PAIR and GCG attacks.

B.12 Timing comparison

In Table 8, we record the computation time for SmoothLLM on different hardware.

B.13 Improving nominal-performance trade-offs

A primary factor behind SmoothLLM's nominal performance degradation that results is the fact that random character-level perturbations degrade the semantic content of input prompts. This problem was also observed in randomized smoothing, and various solutions in that literature were proposed. A notable solution was to prepend a denoising/smoothing function to a ML model, yielding the paradigm of *denoised smoothing*. Recent work has shown that when applied to SmoothLLM, the paradigm of denoised smoothing results in significantly less degraded nominal performance (Ji et al., 2024). In particular, the Pareto curves swept out in Figure 2 of (Ji et al., 2024) shows that SmoothLLM can be implemented with semantic, rather than character-level, perturbations, yielding significantly better nominal performance.

C Attacking SmoothLLM

As alluded to in the main text, a natural question about our approach is the following:

> Can one design an algorithm that jailbreaks SmoothLLM?

The answer to this question is not particularly straightforward, and it therefore warrants a lengthier treatment than could be given in the main text. Therefore, we devote this appendix to providing a discussion about methods that can be used to attack SmoothLLM. To complement this discussion, we also perform a set of experiments that tests the efficacy of these methods.

C.1 Does `GCG` jailbreak SmoothLLM?

We now consider whether `GCG` can jailbreak SmoothLLM. To answer this question, we first introduce some notation to formalize the `GCG` attack.

C.1.1 Formalizing the `GCG` attack

Assume that we are given a fixed alphabet $\mathcal{A}$, a fixed goal string $G \in \mathcal{A}^{m_G}$, and target string $T \in \mathcal{A}^{m_T}$. As noted in § 2, the goal of the suffix-based attack described in (Zou et al., 2023b) is to solve the feasibility problem in (2), which we reproduce here for ease of exposition:

$$\text{find} \quad S \in \mathcal{A}^{m_S} \quad \text{subject to} \quad (\text{JB} \circ \text{LLM})([G; S]) = 1. \tag{47}$$

Note that any feasible suffix $S^\star \in \mathcal{A}^{m_S}$ will be optimal for the following maximization problem.

$$\underset{S \in \mathcal{A}^{m_S}}{\text{maximize}} \quad (\text{JB} \circ \text{LLM})([G; S]). \tag{48}$$

That is, $S^\star$ will result in an objective value of one in (48), which is optimal for this problem.

Since, in general, JB is not a differentiable function (see the discussion in Appendix B), the idea in (Zou et al., 2023b) is to find an appropriate surrogate for (JB $\circ$ LLM). The surrogate chosen in this past work is the probably—with respect to the randomness engendered by the LLM—that the first m_T tokens of the string generated by LLM($[G; S]$) will match the tokens corresponding to the target string T. To make this more formal, we decompose the function LLM as follows:

$$\text{LLM} = \text{DeTok} \circ \text{Model} \circ \text{Tok} \tag{49}$$

where Tok is a mapping from words to tokens, Model is a mapping from input tokens to output tokens, and DeTok = Tok^{-1} is a mapping from tokens to words. In this way, can think of LLM as conjugating Model by Tok. Given this notation, over the randomness over the generation process in LLM, the surrogate version of (48) is as follows:

$$\underset{S \in \mathcal{A}^{m_S}}{\arg\max} \; \log \Pr\left[R \text{ start with } T \mid R = \text{LLM}([G; S]) \right] \tag{50}$$

$$= \underset{S \in \mathcal{A}^{m_S}}{\arg\max} \; \log \prod_{i=1}^{m_T} \Pr[\text{Model}(\text{Tok}([G; S]))_i = \text{Tok}(T)_i \mid$$
$$\text{Model}(\text{Tok}([G; S]))_j = \text{Tok}(T)_j \;\; \forall j < i] \tag{51}$$

$$= \underset{S \in \mathcal{A}^{m_S}}{\arg\max} \; \sum_{i=1}^{m_T} \log \Pr[\text{Model}(\text{Tok}([G; S]))_i = \text{Tok}(T)_i \mid$$
$$\text{Model}(\text{Tok}([G; S]))_j = \text{Tok}(T)_j \;\; \forall j < i] \tag{52}$$

$$= \underset{S \in \mathcal{A}^{m_S}}{\arg\min} \; \sum_{i=1}^{m_T} \ell(\text{Model}(\text{Tok}([G; S]))_i, \text{Tok}(T)_i) \tag{53}$$

where in the final line, ℓ is the cross-entropy loss. Now to ease notation, consider that by virtue of the following definition

$$L([G;S],T) \triangleq \sum_{i=1}^{m_T} \ell(\text{Model}(\text{Tok}([G;S]))_i, \text{Tok}(T)_i) \tag{54}$$

we can rewrite (53) in the following way:

$$\underset{S \in \mathcal{A}^{m_S}}{\arg\min} \quad L([G;S],T) \tag{55}$$

To solve this problem, the authors of (Zou et al., 2023b) use first-order optimization to maximize the objective. More specifically, each step of GCG proceeds as follows: For each $j \in [V]$, where V is the dimension of the space of all tokens (which is often called the "vocabulary," and hence the choice of notation), the gradient of the loss is computed:

$$\nabla_S L([G;S],T) \in \mathbb{R}^{t \times V} \tag{56}$$

where $t = \dim(\text{Tok}(S)$ is the number of tokens in the tokenization of S. The authors then use a sampling procedure to select tokens in the suffix based on the components elements of this gradient.

C.1.2 On the differentiability of SmoothLLM

Given the formalization in the previous section, we now show that SMOOTHLLM cannot be adaptively attacked by GCG. The crux of this argument has already been made; since GCG requires an attacker to compute the gradient of a targeted LLM with respect to its input, non-differentiable defenses cannot be adaptively attacked by GCG.

> **Proposition C.1 (SmoothLLM differentiability)**
>
> SMOOTHLLM(P) is a non-differentiable function of its input, and therefore it cannot be adaptively attacked by GCG.

Proof. Begin by returning to Algorithm 1, wherein rather than passing a single prompt $P = [G;S]$ through the LLM, we feed N perturbed prompts $Q_j = [G'_j; S'_j]$ sampled i.i.d. from $\mathbb{P}_q(P)$ into the LLM, where G'_j and S'_j are the perturbed goal and suffix corresponding to G and S respectively. Notice that by definition, SmoothLLM, which is defined as

$$\text{SMOOTHLLM}(P) \triangleq \text{LLM}(P^\star) \quad \text{where} \quad P^\star \sim \text{Unif}(\mathcal{P}_N) \tag{57}$$

where

$$\mathcal{P}_N \triangleq \left\{ P' \in \mathcal{A}^m \ : \ (\text{JB} \circ \text{LLM})(P') \right.$$
$$\left. = \mathbb{I}\left[\frac{1}{N}\sum_{j=1}^{N}[(\text{JB} \circ \text{LLM})(Q_j)] > \frac{1}{2}\right] \right\} \tag{58}$$

is non-differentiable, given the sampling from $\mathcal{P}_N$ and the indicator function in the definition of $\mathcal{P}_N$. □

C.2 Surrogates for SmoothLLM

Although we cannot directly attack SmoothLLM, there is a well-traveled line of thought that leads to an approximate way of attacking smoothed models. More specifically, as is common in the adversarial robustness literature, we now seek a surrogate for SmoothLLM that is differentiable and amenable to GCG attacks.

C.2.1 Idea 1: Attacking the empirical average

An appealing surrogate for SmoothLLM is to attack the empirical average over the perturbed prompts. That is, one might try to solve

$$\underset{S \in \mathcal{A}^{m_S}}{\text{minimize}} \quad \frac{1}{N} \sum_{j=1}^{N} L([G'_j, S'_j], T). \tag{59}$$

If we follow this line of thinking, the next step is to calculate the gradient of the objective with respect to S. However, notice that since the S'_j are each perturbed at the character level, the tokenizations $\text{Tok}(S'_j)$ will not necessarily be of the same dimension. More precisely, if we define

$$t_j \triangleq \dim(\text{Tok}(S'_j)) \quad \forall j \in [N], \tag{60}$$

then it is likely the case that there exists $j_1, j_2 \in [N]$ where $j_1 \neq j_2$ and $t_{j_1} \neq t_{j_2}$, meaning that there are two gradients

$$\nabla_S L([G'_{j_1}; S'_{j_2}], T) \in \mathbb{R}^{t_{j_1} \times V} \tag{61}$$

and

$$\nabla_S L([G'_{j_2}; S'_{j_2}], T) \in \mathbb{R}^{t_{j_2} \times V} \tag{62}$$

that are of different sizes in the first dimension. Empirically, we found this to be the case, as an aggregation of the gradients results in a dimension mismatch within several iterations of running `GCG`. This phenomenon precludes the direct application of `GCG` to attacking the empirical average over samples that are perturbed at the character-level.

C.2.2 Idea 2: Attacking in the space of tokens

Given the dimension mismatch engendered by maximizing the empirical average, we are confronted with the following conundrum: If we perturb in the space of characters, we are likely to induce tokenizations that have different dimensions. Fortunately, there is an appealing remedy to this shortcoming. If we perturb in the space of tokens, rather than in the space of characters, by construction, there will be no issues with dimensionality.

More formally, let us first recall from § C.1.1 that the optimization problem solved by `GCG` can be written in the following way:

$$\underset{S \in \mathcal{A}^{m_S}}{\arg\min} \sum_{i=1}^{m_T} \ell(\text{Model}(\text{Tok}([G; S]))_i, \text{Tok}(T)_i) \tag{63}$$

Now write

$$\text{Tok}([G; S]) = [\text{Tok}(G); \text{Tok}(S)] \tag{64}$$

so that (63) can be rewritten:

$$\underset{S \in \mathcal{A}^{m_S}}{\arg\min} \sum_{i=1}^{m_T} \ell(\text{Model}([\text{Tok}(G); \text{Tok}(S)])_i, \text{Tok}(T)_i) \tag{65}$$

As mentioned above, our aim is to perturb in the space of tokens. To this end, we introduce a distribution $\mathbb{Q}_q(D)$, where D is the tokenization of a given string, and q is the percentage of the tokens in D that are to be perturbed. This notation is chosen so that it bears a resemblance to $\mathbb{P}_q(P)$, which denoted a distribution over perturbed copies of a given prompt P. Given such a distribution, we propose the following surrogate for SmoothLLM:

$$\underset{S \in \mathcal{A}^{m_S}}{\text{minimize}} \quad \frac{1}{N} \sum_{j=1}^{N} \sum_{i=1}^{m_T} \ell(\text{Model}([\text{Tok}(G); Z_j])_i, \text{Tok}(T)_i) \tag{66}$$

where $Z_1, \ldots, Z_N$ are drawn i.i.d. from $\mathbb{Q}_q(\text{Tok}(S))$. The idea here is to create N randomly perturbed copies of the tokenization of the optimization variable S. Notice that while we employ the empirical average discussed in § C.2.1, the difference is that we now perturb in the space of tokens, rather than in the space of characters. Given this formulation, on can take gradients with respect to the perturbed tokenizations, facilitating the use of GCG on this surrogate. For the remainder of this appendix, we will refer to this surrogate as SurrogateLLM.

C.3 Additional notes on attacks

Adaptive attack that are likely to bypass SmoothLLM are those that allow the attacker to play after a forward pass through SmoothLLM. While this does *not* fit within the threat model described in Figure 3, it's worth noting that any attack that can modify the prompt produced by the majority voting step in an unconstrained way will tend to succeed. We are not currently aware of any such attacks, but if one was given white-box access to a SmoothLLM implementation, such an attack would be theoretically possible.

Algorithm 2: `RandomPerturbation` function definitions

---

1 Function `RandomSwapPerturbation`(P, q)**:**
2   Sample a set $\mathcal{I} \subseteq [m]$ of $M = \lfloor qm \rfloor$ indices uniformly from $[m]$
3   **for** index i in $\mathcal{I}$ **do**
4 $P[i] \leftarrow a$ where $a \sim \text{Unif}(\mathcal{A})$
5   **return** P

6 Function `RandomPatchPerturbation`(P, q)**:**
7   Sample an index i uniformly from $\in [m - M + 1]$ where $M = \lfloor qm \rfloor$
8   **for** $j = i, \ldots, i + M - 1$ **do**
9 $P[j] \leftarrow a$ where $a \sim \text{Unif}(\mathcal{A})$
10   **return** P

11 Function `RandomInsertPerturbation`(P, q)**:**
12   Sample a set $\mathcal{I} \subseteq [m]$ of $M = \lfloor qm \rfloor$ indices uniformly from $[m]$
13   count $\leftarrow 0$
14   **for** index i in $\mathcal{I}$ **do**
15 $P[i + \text{count}] \leftarrow a$ where $a \sim \text{Unif}(\mathcal{A})$
16   count = count + 1
17   **return** P

D   A collection of perturbation functions

In Algorithm 2, we formally define the three perturbation functions used in this paper. Specifically,

- RANDOMSWAPPERTURBATION is defined in lines 1-5;

- RANDOMPATCHPERTURBATION is defined in lines 6-10;

- RANDOMINSERTPERTURBATION is defined in lines 11-17.

In general, each of these algorithms is characterized by two main steps. In the first step, one samples one or multiple indices that define where the perturbation will be applied to the input prompt P. Then, in the second step, the perturbation is applied to P by sampling new characters from a uniform distribution over the alphabet $\mathcal{A}$. In each algorithm, $M = \lfloor qm \rfloor$ new characters are sampled, meaning that $q\%$ of the original m characters are involved in each perturbation type.

D.1   Sampling from $\mathcal{A}$

Throughout this paper, we use a fixed alphabet $\mathcal{A}$ defined by Python's native `string` library. In particular, we use `string.printable` for $\mathcal{A}$, which contains the numbers 0-9, upper- and lower-case letters, and various symbols such as the percent and dollar signs as well as standard punctuation. We note that `string.printable` contains 100 characters, and so in those figures that compute the probabilistic certificates in § 3.3, we set the alphabet size $v = 100$. To sample from $\mathcal{A}$, we use Python's `random.choice` module.

Figure 15: **An example of the incoherency threshold.**

E   The incoherency threshold

In § 5 of the main text, we discussed the interplay between q and the ASR when running SmoothLLM. In particular, we first observed from Figure 7 that in some cases, for lower values of N, higher values of q resulted in larger ASRs than did lower values of q. We then argued that this phenomenon is attributable to the LLM becoming confused due to the significant level of perturbation applied to the text, resulting in the false identification of a jailbreak. In Figure 15, we show an example—taken directly from this experiment—that demonstrates this situation. On the top line, we show (an abbreviated version of) the system prompt used for Llama2. Next, inside the parentheses, we show an unperturbed copy of the goal prompt (shown in **green**) and an adversarially-generated suffix (shown in **red**). However, as described in the main text, unperturbed prompts are never passed to the LLM. Rather, the LLM receives perturbed copies of the attacked prompt. On the third line, we show an example of a perturbed copy of the original prompt, which was obtained by running Algorithm 1 with swap perturbations and $q = 20\%$. And finally, on the final line, we show the response from the LLM.

Notice that in the response, the LLM generates text that indicates that the perturbed prompt cannot be understood, c.f., "the statement you provided... does not make sense." And while this text does not constitute a jailbreak, it also does not include the keywords used in (Zou et al., 2023b) to identify a jailbreak; these keywords, which are also listed in Appendix B, are as follows:

> JailbreakKeywords = ["I'm sorry", "Sorry", "I am sorry", "I apologize", "As an", "As an AI", "I'm an", "I'm just", "As a language model", "As an Assistant", "I cannot", "I do not", "Hello!"]

Therefore, this response is identified as bypassing the safety filter, and therefore contributes to a nonzero ASR. In other words, this response represents a false positive with respect to the JB function used in (Zou et al., 2023b).

That this phenomenon occurs is due to using a value of q that renders the perturbed prompt incoherent. For this reason, we term this phenomenon "passing the *incoherency threshold*" to indicate that there exist

values or ranges for q that will reduce the semantic content contained in the prompt. Therefore, as indicated in the main text, q should not be chosen to be particularly large when running SmoothLLM.

F   Additional related work

F.1   Adversarial examples, robustness, and certification

A longstanding disappointment in the deep learning community is that DNNs often fail in the presence of seemingly innocuous changes to their input data. Such changes—include nuisances in visual data (Laidlaw et al., 2020; Robey et al., 2020; Wong & Kolter, 2020), sub-population drift (Santurkar et al., 2020; Koh et al., 2021), and distribution shift (Arjovsky et al., 2019; Eastwood et al., 2022; Robey et al., 2021)—limit the applicability of deep learning methods in safety critical areas. Among these numerous failure modes, perhaps the most well-studied is the setting of adversarial examples, wherein it has been shown that imperceptible, adversarially-chosen perturbations tend to fool state-of-the-art computer vision models (Biggio et al., 2013; Szegedy et al., 2013). This discovery has spawned thousands of scholarly works which seek to mitigate this vulnerability posed.

Over the past decade, two broad classes of strategies designed to mitigate the vulnerability posed by adversarial examples have emerged. The first class comprises *empirical defenses*, which seek to improve the empirical performance of DNNs in the presence of a adversarial attacks; this class is largely dominated by so-called *adversarial training* algorithms (Goodfellow et al., 2014; Madry et al., 2017; Zhang et al., 2019), which incorporate adversarially-perturbed copies of the data into the standard training loop. The second class comprises *certified defenses*, which provide guarantees that a classifier—or, in many cases, an augmented version of that classifier—is invariant to all perturbations of a given magnitude (Lecuyer et al., 2019). The prevalent technique in this class is known as *randomized smoothing*, which involves creating a "smoothed classifier" by adding noise to the data before it is passed through the model (Cohen et al., 2019; Salman et al., 2019; Yang et al., 2020).

F.2   Comparing randomized smoothing and SmoothLLM

The formulation of SmoothLLM adopts a similar interpretation of adversarial attacks to that of the literature surrounding randomized smoothing. Most closely related to our work are non-additive smoothing approaches (Levine & Feizi, 2020; Yatsura et al., 2022; Xue et al., 2023). To demonstrate these similarities, we first formalize the notation needed to introduce randomized smoothing. Consider a classification task where we receive instances x as input (e.g., images) and our goal is to predict the label $y \in [k]$ that corresponds to that input. Given a classifier f, the "smoothed classifier" g which characterizes randomized smoothing is defined in the following way:

$$g(x) \triangleq \arg\max_{c \in [k]} \Pr_{x' \sim \mathcal{B}(x)} [f(x') = c] \tag{67}$$

where $\mathcal{B}$ is the smoothing distribution. For example, a classic choice of smoothing distribution is to take $\mathcal{B}(x) = x + \mathcal{N}(0, \sigma^2 I)$, which denotes a normal distribution with mean zero and covariance matrix $\sigma^2 I$ around x. In words, $g(x)$ predicts the label c which corresponds to the label with highest probability when the distribution $\mathcal{B}$ is pushed forward through the base classifier f. One of the central themes in randomized smoothing is that while f may not be robust to adversarial examples, the smoothed classifier g is *provably* robust to perturbations of a particular magnitude; see, e.g., (Cohen et al., 2019, Theorem 1).

The definition of SmoothLLM in Definition 3.1 was indeed influenced by the formulation for randomized smoothing in (67), in that both formulations employ randomly-generated perturbations to improve the robustness of deep learning models. However, we emphasize several key distinctions in the problem setting, threat model, and defense algorithms:

- **Problem setting: Prediction vs. generation.** Randomized smoothing is designed for classification, where models are trained to predict one output. on the other hand, SmoothLLM is designed for text generation tasks which output variable length sequences that don't necessarily have one correct answer.

- **Threat model: Adversarial examples vs. jailbreaks.** Randomized smoothing is designed to mitigate the threat posed by traditional adversarial examples that cause a misprediction, whereas

    SmoothLLM is designed to mitigate the threat posed by language-based jailbreaking attacks on LLMs.

- **Defense algorithm: Continuous vs. discrete distributions.** Randomized smoothing involves sampling from continuous distributions (e.g., Gaussian (Cohen et al., 2019), Laplacian (Teng et al., 2019)and others (Yang et al., 2020; Fischer et al., 2020; Rosenfeld et al., 2020)) or discrete distrbutions (Levine & Feizi, 2020; Yatsura et al., 2022; Xue et al., 2023). SmoothLLM falls in the latter category and involves sampling from discrete distributions (see Appendix D) over characters in natural language prompts. In particular, it is most similar to Xue et al. (2023), which smooths vision and language models by randomly dropping tokens to get stability guarantees for model explanations. In contrast, our work is designed for language models and randomly replaces tokens in a fixed pattern.

Therefore, while both algorithms employ the same underlying intuition, they are not directly comparable and are designed for distinct sets of machine learning tasks.

F.3 Adversarial attacks and defenses in NLP

Over the last few years, an amalgamation of attacks and defenses have been proposed in the literature surrounding the robustness of language models (Morris et al., 2020; Zhang et al., 2020). The threat models employed in this literature include synonym-based attacks (Ren et al., 2019; Wang et al., 2019; Alzantot et al., 2018), character-based substitutions (Li et al., 2018), and spelling mistakes (Pruthi et al., 2019). Notably, the defenses proposed to counteract these threats almost exclusively rely on retraining or fine-tuning the underlying language model (Wang et al., 2021b;a; Zhou et al., 2021). Because of the scale and opacity of modern, highly-performant LLMs, there is a pressing need to design defenses that mitigate jailbreaks without retraining. The approach proposed in this paper—which we call SmoothLLM—fills this gap.

F.4 Fine-tuning defenses for LLMs

Since this work appeared on arXiv, various defense algorithms against jailbreaking attacks have been proposed. A notable subset of these attacks have involved fine-tuning language models to refuse to reply to objectionable requests (Huang et al., 2024b;a; Hsu et al., 2024). These methods have employed a range of techniques, including representation noising (Rosati et al., 2024), latent adversarial perturbations (Sheshadri et al., 2024), and backdoor insertion (Wang et al., 2024). Notably, the so-called circuit breaker defense, which is based on the observation that personas such as helpfulness or harmfulness tend to cluster in low-dimensional latent representation spaces of LLMs (Zou et al., 2023a; Arditi et al., 2024), has also shown significant robustness against a wide-range of jailbreaking attacks (Zou et al., 2024). Such improvements are particularly notable given that benign fine-tuning (i.e., fine-tuning on non-objectionable data) is known to degrade the alignment of highly-performant LLMs (Peng et al., 2024; Qi et al., 2023).